


AFRL-IF-RS-TR-1999-117
Final Technical Report
June 1999



ASSESSMENT OF 2-PHOTON OPTICAL STORAGE POTENTIALS FOR THE YEARS 2000 - 2010

Call/Recall, Inc.

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. A28E

19990719 121

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-1999-117 has been reviewed and is approved for publication.

APPROVED:



ALBERT A. JAMBERDINO
Project Engineer

FOR THE DIRECTOR:



JOHN V. MCNAMARA, Technical Advisor
Information & Intelligence Exploitation Division
Information Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/IFED, 32 Brooks Road, Rome, NY 13441-4114. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

**ASSESSMENT OF 2-PHOTON OPTICAL STORAGE POTENTIALS
FOR THE YEARS 2000 - 2010**

F. B. McCormick
Sadik Esener

Contractor: Call/Recall, Inc.
Contract Number: F30602-97-C-0218
Effective Date of Contract: 31 May 1997
Contract Expiration Date: 30 June 1999
Short Title of Work: Assessment of 2-Photon Optical Storage
Potentials for the Years 2000 - 2010

Period of Work Covered: May 97 - Jun 99

Principal Investigator: F. B. McCormick
Phone: (619) 550-0596
AFRL Project Engineer: Albert A. Jamberdino
Phone: (315) 330-2845

Approved for public release; distribution unlimited.

This research was supported by the Defense Advanced Research
Projects Agency of the Department of Defense and was monitored by
Albert A. Jamberdino, AFRL/IFED, 32 Brooks Road, Rome, NY
13441-4114.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1999		3. REPORT TYPE AND DATES COVERED Final May 97 - Jun 99
4. TITLE AND SUBTITLE ASSESSMENT OF 2-PHOTON OPTICAL STORAGE POTENTIALS FOR THE YEARS 2000 - 2010			5. FUNDING NUMBERS C - F30602-97-C-0218 PE - 62712E PR - A28E0000 TA - 00 WU - 00	
6. AUTHOR(S) F. B. McCormick and Sadik Esener				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Call/Recall, Inc. 6160 Lusk Blvd, C-206 San Diego CA 92121			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency Air Force Research Laboratory/IFED 3701 North Fairfax Drive 32 Brooks Road Arlington VA 22203-1714 Rome NY 13441-4114			10. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-1999-117	
11. SUPPLEMENTARY NOTES Air Force Research Laboratory Project Engineer: Albert A. Jamberdino/IFED/(315) 330-2845				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This program includes an assessment of Government and Industrial applications for optical data storage, and identifies the windows of opportunity by determining current state of the art, as well as future advanced technology predictions. This study included investigations into magnetic, holographic, probe recording, 2 photon, 3-D memories, etc. The program culminated in a workshop describing the quantified results.				
14. SUBJECT TERMS Photonics, Optical Photochromic Materials, 3-D Memory			15. NUMBER OF PAGES 98	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

1. Executive Summary	1
2. Scope and Methods	2
3. Military and Commercial Applications	3
3.1. Commercial applications and markets	3
3.1.1. Internet impact on data storage	5
3.1.2. Multimedia authoring and network servers	7
3.1.3. Data warehousing and mining	7
3.1.4. Niche applications	9
3.2. Military applications	9
3.2.1. Surveillance and reconnaissance: the UAV evolution	9
3.2.2. Training and simulation: synthetic environments	12
3.2.3. Mission support: Transportable Training and "World Map"	14
3.3. Potential synergies between military and commercial roadmaps	18
4. Conventional Storage Technology	21
4.1. Optical disk hardware review	21
4.1.1. Phase change media projections	24
4.1.2. Magneto-optical media projections	25
4.2. Channel Electronics and Peripheral Interconnection Trends	25
4.2.1. Channel Electronics	26
4.2.2. Interconnection Trends	31
4.3. Magnetic recording trends	38
5. Advanced Optical Storage Technologies	40
5.1. Parallel Access	40
5.2. Near-field optical recording (Terastor)	44
5.3. Volume holography	46
5.4. Persistent spectral hole burning / temporal holography	55
5.5. Pit Depth Modulation	57
5.6. Probe storage	58
6. 2-Photon Optical Storage Projections	62
6.1. Overview and progress	62
6.2. Projections	65
6.2.1. Recording	65
6.2.2. Reading: the FROST architecture	67
6.2.3. FROST performance and scaling projections	68
7. Windows of opportunity	71
7.1. Windows of opportunities for Two-Photon Disk Storage	72

7.1.1.	Potential product evolution path	74
8.	Appendix I: Projection Validation Workshop	75
8.1.	Workshop agenda	76
8.2.	Workshop attendees	77
9.	References	78

List of Figures

Figure 1.	Potential evolution of application requirements for removable storage.....	5
Figure 2.	Potential commercial applications for optical storage systems: projected data capacity requirements.....	6
Figure 3.	Potential commercial applications for optical storage systems: projected data transfer rate requirements.....	6
Figure 4.	Potential commercial applications for optical storage systems: projected data access time requirements.....	7
Figure 5.	Range of M&S Embraced by the DoD M&S Vision.....	13
Figure 6.	Goddard Earth Observing System (EOS) Data Assimilation System 1998 hardware requirements and projections. (from reference 30).....	14
Figure 7.	Potential military applications for optical storage systems: projected data capacity requirements.....	15
Figure 8.	Potential military applications for optical storage systems: projected data rate requirements.....	16
Figure 9.	Potential military applications for optical storage systems: projected data access time requirements.....	16
Figure 10.	Potential overlap of mass-market and military applications.	18
Figure 11.	Collaborative telepresence system concept.	19
Figure 12.	1995 economic impact of optoelectronics (adapted from OIDA).	19
Figure 13.	Conventional optical disk pickup head architecture.....	21
Figure 14.	Functional diagram of an optical head.....	21
Figure 15.	Performance comparisons of conventional data storage technologies (1997).	23
Figure 16.	Example roadmap describing performance potentials of phase change media.....	24
Figure 17.	Example roadmap describing performance potentials of MO media based systems ..	25
Figure 18:	Read channel functional blocks for disk based storage devices.....	26
Figure 19:	Comparison of read channel transfer rates with peripheral interconnect (PI) technology. Commodity PI follows the trend in desktop interconnects, such as SCSI and Firewire. High performance PI follows the HIPPI standards for interconnecting workstation/supercomputer clusters.....	27
Figure 20:	Projected internal data rates for various storage technologies. Internal data rate is derived by multiplying the linear bit density by the rotational velocity of the disk.	28
Figure 21.	Channel electronics roadmap for serial read out.	30

Figure 22. Network-to-storage system bottleneck: performance requirements for various network-attached services (from reference 37).....	31
Figure 23: Aggregate throughput for interconnect technologies applicable to storage devices. Data points represent commercial products. The graph highlights the trends in high performance versus commodity interconnects.....	36
Figure 24: The number of data channels used for various PI interconnects versus the year of implementation.	37
Figure 25. Gap between peripheral interconnection rates and serial read channel rates (high performance is an extension of HIPPI, commodity is an extension of SCSI, serial read channel is an extension of PRML implementations).	37
Figure 26. Multi-beam optical head investigated at Fujitsu.....	40
Figure 27. Kenwood's \$99 MULTI BEAM 40X Plus™ CD-ROM drive (source: http://www.kenwoodtech.com/home_toplink.html)	41
Figure 28. Sustained data transfer rate comparisons for serial CD and DVD –ROM and 6-channel parallel CD-ROM drives (source: http://www.zenresearch.com/product_info4.html).	41
Figure 29. Multibeam optical head schematic for Zen CD-ROM drives (source: http://www.zenresearch.com/product_info4.html).	42
Figure 30. Schematic of a VCSEL array access with a massively parallel near field optics architecture to optical disks promising TB capacities and Tb/s data transfer.	42
Figure 31. Use of a VCSEL for recording and readout in a near field optics geometry.....	43
Figure 32. Terastor's near-field optical approach using MO media, and their projected areal density roadmap comparison.	45
Figure 33. Principles of 90 degree geometry holographic recording and reading.....	46
Figure 34. Holographic demonstration systems constructed at Rockwell Sciences Center.	50
Figure 35. Temporal holography using persistent spectral hole burning.....	57
Figure 36. Pit depth modulation: principle and result.....	58
Figure 37. Probe storage at Matsushita: Record breaking erasable spots of 10 nm diameter have been recorded in phase change media using STM.....	60
Figure 38. Thermo-mechanical recording at IBM using rotating plastic disk using an atomic force microscope tip. At selected locations, electrical current was pulsed through the tip causing the tip to be heated. The heat from the tip softens the substrate and the pressure from the tip.....	61
Figure 39. Principle of 2-photon data recording and readout.	62
Figure 40. Orthogonally addressed 2-photon recording architectures.	63
Figure 41. Portable ROM reader: (a) schematic, (b) portable reader.	63
Figure 42. Example image, data plane, and side view of the 100 layer, 8mm thick cube.....	64
Figure 43. (a) Automated disk recorder, and (b) disk test stand.....	64
Figure 44. (a) Colinearly illuminated disk and data signal, and (b) side view of 17 layer 2-photon 3D disk.	65
Figure 45. Preliminary recording results achieved using a 76 MHz, 200 fs pulse length near infrared laser.	66
Figure 46. Fast Readout Optical Storage Technology (FROST) approach.....	67

Figure 47. Target areal densities for mass-market advanced optical storage technologies. (Source: Marvin Keshner, HP Labs, in the High Density Optical Data Storage Workshop, San Diego, CA, March 2, 1998.).....	70
Figure 48. Possible road map for alternative technologies	73
Figure 49. Potential evolution of capacity vs. transfer rate and comparison with conventional techniques (Based Call/Recall 1997 storage survey).....	73

1. Executive Summary

The dawn of the multimedia age has accelerated the need for high-capacity, high-performance data storage systems for evolving applications in both military and commercial arenas. The well-known progress in processor performance has to this point been matched by progress in peripheral interconnection links and data storage devices. However, while the paths for continued progress in processors and interconnect are clearly marked, and potential near-term solutions for data storage are being developed, the path for continued progress after 3-5 years is less clear. This purpose of this study is to clarify the path for optical data storage in the next decade, and determine the role of 2-photon optical data storage within this vision. Based on DARPA feedback from the kickoff of this program, this study seeks to provide:

- A compelling reason for high-capacity, fast-transfer-rate storage in DoD applications.
- Detailed application overviews and assessments to define metrics for novel scalable storage systems.
- The characteristics that will be the needed in a revolutionary storage technology in 10 years.
- Comparisons with current and projected competing techniques.
- Reasons why commercial markets will support the development of future components and systems.
- Estimates of when potential technologies and the associated user-base be ready.

To accomplish this, commercial applications currently driving optical storage progress have been analyzed, key military applications have been investigated, and the evolving military application requirements have been estimated. Together, these analyses show the future directions and application roadmaps. Conventional optical storage approaches and their associated commercial roadmaps have been surveyed to determine the critical issues within the roadmaps, especially concerning data transfer rates from the media and to/from the host or network. Promising advanced storage technologies have been investigated to determine their current status and potential, including the status and roadmap for 2-photon storage technology. To enable an optical data storage technology to succeed in recapturing some market share for the US, and further stimulate the US optoelectronics industry, the potential windows of opportunity and associated performance requirements have been determined.

2. Scope and Methods

The purpose of this study is to assess the potential of 2-photon storage techniques in satisfying the data storage needs of future military applications. The approach used to answer this question was to identify and project the storage requirements for critical military applications in the next decade, and, by projecting the progress of conventional and novel optical storage technologies (e.g. - 2-photon and others), determine the ability of optical storage to meet these needs.

This study was conducted by Rick McCormick and Sadik Esener at Call/Recall, Inc. (Prime Contractor), Charlie Kuznia and Alexander Sawchuk at the University of Southern California (commercial applications, channel projections, and peripheral interconnection projections), and Susan Hunter and Waguhi Ishak at Hewlett-Packard Company (commercial application projections). The methods used in this study may be grouped into 3 areas: traditional and Internet literature surveys, technical meeting and workshop attendance and information gathering, and personal and telephone interviews with individuals and companies leading various aspects of optical storage technology and application. Key information gathering opportunities were meetings such as Optical Data Storage '97 and '98, the PSHB '98 workshop, SPIE symposium on Advanced Optical Memories and Interfaces to Computer Systems '98. Additionally, the study also includes input gathered from High Density Optical Data Storage Workshop (organized and conducted under this contract), and from the WTEC panel study of High Density Data Storage in Japan, led by one of the PI's (Esener) on this study.

The results of this study are presented in the following format. Critical commercial and military applications are identified and their requirements discussed in section 3. Section 4 projects the progress of conventional optical storage technologies, peripheral interconnection trends, and magnetic disk storage. The differences in the application requirements and the storage roadmaps point to potential windows of opportunity for new storage technologies, and these windows are identified in section 5. The principles used, present status and roadmap of 2-photon optical storage is discussed in section 6, and future projections for other novel storage techniques are presented in section 7. The agenda and results of the High Density Optical Data Storage Workshop conducted under this contract are summarized in section 8.

3. Military and Commercial Applications

As society progresses from the information age to the multimedia age, our storage requirements have similarly progressed from the 1.44 MB magnetic floppy disks storing text to CD-recordable media, Zip-disks, and, in Asia, magneto-optic mini disks storing text, images, and video. This trend will likely accelerate in the next decade, as image and video sensing platforms multiply, DVD movies become ubiquitous, video teleconferencing becomes commonplace, and virtual reality education and entertainment applications enter the mainstream. All of these coming applications share the need for high capacity and high data transfer rate. Additionally, they often require high reliability, and removability for security or distribution reasons. As discussed in the following sections, projected military and commercial application requirements both show similar performance requirements, largely due to the similarity of the data being stored (e.g.-imagery and video).

3.1. Commercial applications and markets

Optical storage has typically offered a reliable and removable storage medium with excellent robustness and archival lifetime at very low cost. A key difference of optical recording with respect to magnetic recording is the ease with which the media can be made removable. Both optical recording and readout can be performed with a head positioned relatively far away from the storage medium, unlike magnetic hard drive heads. This allows the medium to be removable and effectively eliminates head crashes, increases reliability. In addition, during recording, optical radiation is used as a focused thermal source allowing the use of more stable materials suitable for archival lifetimes. On the other hand, the remote optical head is heavier and leads to slower access times when compared to hard disk drives.

Consequently, optical storage has remained limited to market segments requiring removability and reliability that are not well served by magnetic hard disks. Typical applications involve archival storage, including software distribution, storing digital photographs and medical imaging, information appliances including recording movies, other video materials, and multimedia presentations at home and business, and on-line databases including video servers. More recently the magnetic tape market for video camcorders and VCRs has also been targeted. These type of applications, while benefiting from random access capabilities of disk systems, are less sensitive to access time requirements but require low cost and high capacity removable storage. As discussed in this report, many emerging military and commercial applications share these requirements of high capacity at low cost, reliability, removability, and robustness.

Optical data storage, which once appeared to be a failing technology in the consumer marketplace, has quickly found its way into homes and offices with the multimedia revolution through CD-ROM and magneto-optical disks. Due to their low-cost replication, high capacity, robustness, and removability, CD-ROM systems have become competitive with magnetic floppy disks for applications such as software distribution and home multimedia applications. The need for high capacity, low cost and high transfer rate storage is becoming more and more critical as

our society becomes more information-based, as indicated in the projections shown in Table 1. For consumer applications the most attractive unique features of optical storage systems are their higher capacity per disk, removability, lower cost and long memory persistence for archival applications. In addition to home multimedia some emerging applications such as multimedia file servers (1 TB), high-definition television and video disk recorder (HDTV/VDR) (100 GB), and “data warehousing” applications (10 petabytes, with 1 GB/s I/O rate), may best be addressed in the near future by optical storage. For most of these applications high capacity and lowest cost per bit are the most critical figure of merits. For these applications, the access time is less critical making optical storage the technology of choice. It is expected that the new two layer DVD systems will further enhance the commercial acceptance of optical data storage. Also, recent joint ventures (e.g.- ASMO) have been formed to further increase the capacity of magneto-optical devices to about 10GB in response to projected explosive markets.

Table 1. Market Demand Forecasts (in millions of units)

	1998	2001
CD systems	90	130
DVD systems	5	45
Video cassette recorders	20	25
Total	115	200

(source: *Wall Street Journal*)

In contrast to non-removable systems, for removable storage, yearly increases in performance are not necessarily desirable. This is because removable storage systems and media are tightly constrained by standards that are established for compatibility purposes. Removable storage manufacturers instead introduce products at an entry capacity and performance point that is desirable for a particular data type. Thus, the optical storage market is essentially driven by applications rather than by progress made in technology. Therefore projections made on optical storage critically depend on application roadmaps such as the one shown in Figure 1. With the DVD ROM standard it is possible to distribute one two-hour video movie per disk. Soon the DVD RAM standard will enable consumers to record a two-hour long video per disk. An important jump in performance will be needed slightly after the turn of the millennium to address HDTV applications requiring 15GB capacities per movie. There are also serious considerations by both the phase change and MO manufacturers to use 30- 40GB disks as VCR replacement, perhaps shortly before the year 2005. Guessing the type of applications that may drive mass-market optical storage technologies beyond 2005 is certainly speculative at this point in time. It is however plausible that 3-D interactive video in some form of virtual reality application might become a driver for higher capacities within the next decade, pushing mass-market capacity requirements beyond 0.1TB to 1TB per disk.




Figure 1. Potential evolution of application requirements for removable storage.

(Source: information gathered from OITDA and surveys performed by Call/Recall Inc. and USC)

Over the next decade with the emergence of the Internet, a new and important class of applications referred to Server-based applications will emerge. While personal applications will define future removable storage standards, server-based applications will significantly boost the size of the optical and magnetic disk storage market. Many Server-based applications like electronic commerce, medicine, and libraries, among others, require modest access times (<10ms) but very large storage capacities and appreciable data transfer rates. These applications, because of their very large capacity requirements will initially be constructed as RAID or disk library systems based on commodity personal computer drives. Drive and database maintenance costs will be among deciding factors for various technology solutions. Figure 2 - Figure 4 summarize some of the data storage requirements of commercial application projections gathered from the OITDA and surveys performed by Call/Recall, Inc. and USC.

3.1.1. Internet impact on data storage¹

Every new Web page that is added requires additional data storage, both for operation on the Web server, and for backups and archiving. The explosive growth in both Web pages and Web users will only mean continued growth in overall data storage. International Data corporation (IDC)² estimates that by the year 2000 there will be 441,000 Web servers shipping per year. Each of these will come configured with significant disk capacity, to say nothing of add-on storage to existing servers. IDC also estimates that there will be 166 million Web users by the year 2000. If each of those download 100 MB of information, that would consume 16.6 quadrillion bytes of storage. That's equivalent to 16.6 million 1 GB disk drives. There is much discussion about inexpensive Internet "terminals" or diskless PCs to access the Internet. In reality, the bandwidth of typical Internet connections is too low to constantly access data across the Intranet. Consequently, key information is downloaded to the disk drive on the PC, and there will be a need for disk storage on these "terminals" until high-speed links are available. In addition, IDC estimates that 4.6 million intranet servers will ship in the year 2000. If each of these comes with 25 GB of capacity, that represents another 115 quadrillion bytes of storage.




Figure 2. Potential commercial applications for optical storage systems: projected data capacity requirements.

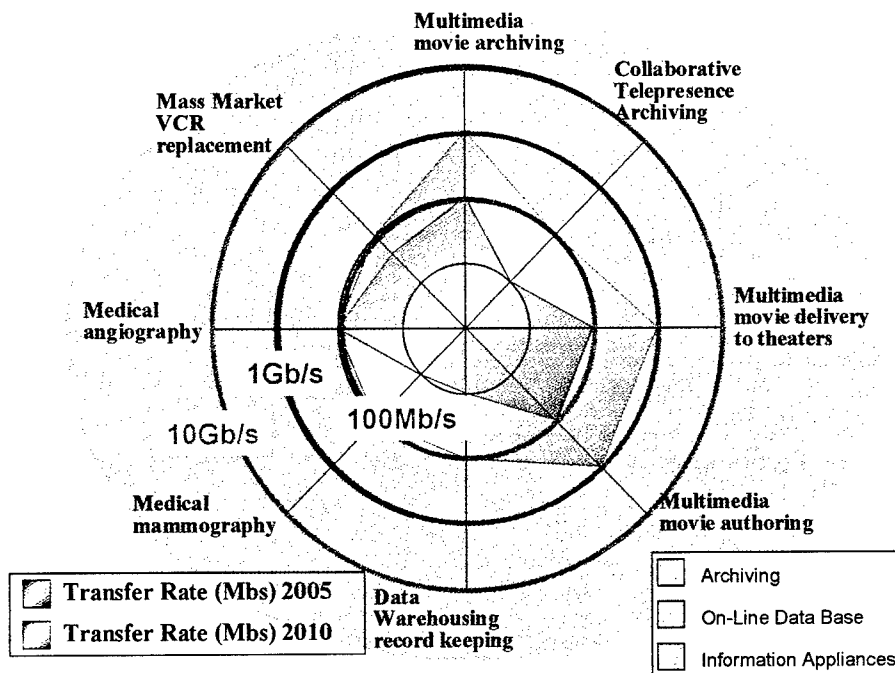


Figure 3. Potential commercial applications for optical storage systems: projected data transfer rate requirements.

3.1.2. Multimedia authoring and network servers

One commercial application with data storage requirements rivaling many military applications is digital movie production. Hollywood is rapidly moving towards digital production technology, both to create special effects, and to mix and edit live action, matte photography, and computer generated effects. One example of this sort of digital production company is Cinesite³, a wholly owned subsidiary of Eastman Kodak Company. Cinesite scans and stores images (movie frames) ranging in size from 12 Mbytes to 100 Mbytes, and can store up to 4 hours of movie footage on-line (>2TB). A typical editing location may use 60 Silicon Graphics servers, linked to over 125 Seagate Barracuda magnetic hard drives in a multiple RAID configuration. To accurately qualify the edited scenes, they must be replayed at HDTV resolution, necessitating a multiple RAID configuration. By striping four RAID systems together, Cinesite has boosted its sustained transfer rates to 106 MB/s, however, still greater storage capacity and data transfer rates are needed.




Figure 4. Potential commercial applications for optical storage systems: projected data access time requirements.

3.1.3. Data warehousing and mining

As digital information become more image and video rich, there is an emerging need in the critical area of "data warehousing" and "data mining". Market driven companies are realizing the key to owning a market segment lies in effectively "mining" their customer information databases in a timely manner. A Meta Group study entitled "Data Mining: Trends, Technology and Implementation Imperatives" predicts that total revenue for this market will hit \$8.4 billion in the next two years, a 150% increase over the \$3.3 billion generated last year.⁴ The driving

factor behind the trend is the move toward customer-centric marketing, which requires specialized data mining applications to better understand consumers' buying patterns and demographic profiles.⁵ First Union Corporation, a \$220 billion bank holding company, recently completed a data warehouse with 4 terabytes in online capacity. The data warehouse software runs across 115 IBM SP2 AIX servers. The data depot consolidates retail information feeds from 20 retail data sources, such as ATMs and credit-card transactions, and stores comprehensive account information on over 16 million consumer banking accounts, over a rolling 24-month window. Virtually every customer transaction resides in the primary data warehouse. Once collected, summary information can be skimmed down to a regional data mart, where it can be quickly sliced and diced by First Union's marketing analysts. Officials of the company say the system will improve sales tracking and trend analysis across 2,400 branches and 3,400 ATMs nationwide to tune the effectiveness of direct-marketing programs and boost bank revenue by some \$100 million annually.

MasterCard International Inc. recently launched a new company that uses MasterCard's huge warehouse of credit-card transaction data to sell information on consumer buying to retailers.⁶ Transactional Data Solutions, a joint venture of MasterCard and market-research firm Symmetrical Resources Inc., plans to offer its first service, Merchant Advisor to provide aggregate consumer buying information based on credit-card transactions and other data without including customer-specific information such as names, addresses, phone numbers, and E-mail addresses. The service is designed to help retailers boost sales and lure customers with more effective advertising and marketing. It compares retailers' performance with competitors over a given period and in different stores. Age, income, buying habits, and other customer data will be provided in clusters, such as women aged 35 to 45 with incomes of \$80,000 to \$100,000. With data from Symmetrical, the reports will also tell retailers which media sources the shoppers tend to see. Merchant Advisor reports will draw on data from MasterCard's 7-terabyte Oracle data warehouse in St. Louis, as well as from other sources. Reports will be provided quarterly for \$75,000 to \$150,000 a year, depending on the size of the retailer. Monthly reports will be available beginning in 2000. National and local versions of the service will be available for apparel, automotive, grocery, home furnishing, and other types of stores.

The increasing multimedia content in many of these databases drives their sizes into the Petabyte range. In this regime, the limited access times (seconds) and transfer rates (10's of Mb/s) of traditional hierarchical storage management (HSM) techniques using optical jukebox or magnetic tape archival severely limit timely access and searching. An emerging solution for this problem is "bulk disk array data warehousing", in which information is kept on large disk arrays in order to improve accessibility and searchability over tape and jukebox archiving. Three key storage system attributes are needed in these systems: protection from data loss, continuous random access to data (on-line), and consistent high data transfer rates. Additionally, given the large number of disks needed in these multi-Petabyte warehouses, individual disk systems in the array must have low system and media costs and high system and data reliability. Thus, there is need in this sort of application for low cost, high capacity, high reliability drives systems. As an example of current state-of-the-art Data Warehousing systems, Compaq's Enterprise Storage

Array 12000 storage system contains as many as 1,000 dual-channel drives/controllers, costs as much as \$1.5 million.⁷ It can store as much as 10 terabytes of data, and can be customized for input/output-intensive mail and messaging applications or for high-bandwidth video streaming. Compaq is also shipping the midrange RAID Array 8000 storage system, which can hold as much as 1.3 terabytes of data; an entry-level 400-Gbyte configuration is priced at \$68,000.

3.1.4. Niche applications

Examples of a niche applications with demanding data storage requirements include data collection for high energy particle physics experiments, and Very-Long-Baseline Interferometry (VLBI).⁸ VLBI has been used to make astronomical images with an angular resolution of better than 100 micro-arc-seconds in the millimeter wavelength bands, a far higher resolution that can presently be achieved with ground-based optical astronomy. VLBI measurements of group delay are also used for long distance geodesy at sub-cm levels of accuracy and provide direct measurements of tectonic plate motion. In both applications, the signal-to-noise ratio (SNR) is proportional to the square root of the number of bits processed, which gives rise to the requirements for high data rate recording. In the 25-year history of VLBI, the data rate has grown by more than three orders of magnitude to the current rate of 1 Gb/s. Despite these gains, VLBI's thirst for bandwidth has not diminished, with planned experiments requiring data transfer rates above 8 Gb/s and capacities greater than 2 TB.

3.2. Military applications

Military victories are now largely determined by who wins the "information war", and military applications projected to have critical data storage needs may be grouped into three broad categories: surveillance/reconnaissance, simulation and training, and mission support. Many surveillance applications are already providing real-time high-resolution imagery, and video, from conventional platforms (aircraft, satellites, etc.) and in the next 10 years, more and more surveillance and reconnaissance data will come from unmanned aerial vehicles (UAVs)⁹. The data storage and processing requirements for real-time surveillance applications in general, and video from UAV and UCAV (unmanned combat aerial vehicles) in particular are projected to grow exponentially over the next 10 years.¹⁰ Additional key military application areas are weapon system trainers and simulators, electronic intelligence data processing, weather predictions, threat system modeling, and war-gaming scenarios. In the following sections, we provide several examples of applications that will drive data storage needs.

3.2.1. Surveillance and reconnaissance: the UAV evolution

Unmanned Aerial Vehicles (UAVs) are expected to play an increasingly dominant role in surveillance and reconnaissance. UAVs are an established part of modern warfare with many countries employing them operationally, primarily for reconnaissance and surveillance but also for decoys, electronic combat, communications relay and, to a limited extent, attack. Countries already using or evaluating the use of UAVs include the United States, Britain, France, Israel, Canada, Japan, South Africa, Singapore, Russia, Poland, Finland, Syria, China, India, Spain,

Morocco, Egypt, Portugal, Italy, Syria, Iraq, Bulgaria, Czech Republic, Argentina, Brazil, Greece, Belgium, Holland, Switzerland, Germany, and Australia.¹¹ Already, UAVs are being used in the war against drugs within the US. UAVs are seeing more and more use in recent international conflicts. RPVs and autonomous UAVs were used by both sides throughout the 1991 Gulf War, primarily as reconnaissance and surveillance platforms. The US, Britain and France made effective use of systems such as Pioneer, Pointer, Exdrone, Midge, Alpylles Mart and the Canadair CL-89, while the Iraqis used the Al Yamamah, Marakeb-1000, Sahreb-1 and Sahreb-2. For Operation Desert Shield and Operation Desert Storm coalition tactical reconnaissance UAVs flew a total of 530 missions, logging 1,700 hours aloft. By providing affordable, real-time or near-real-time intelligence direct to the commander on the spot, they were a key element of the intelligence and reconnaissance systems available to the US-led Coalition Forces. This led Vice Chairman of the US Joint Chiefs of Staff, Admiral William Owens to propose:

*"A notional future 'system of systems' that links reconnaissance satellites and UAVs in surveillance of a 200 nautical miles by 200 nautical mile battle area."*¹²

Post-Gulf War, preventing the loss of human life in combat has become a primary concern, which helped solidify the roles of UAVs in future conflicts. Significant use of UAVs has been with the UN peacekeeping effort in the former Yugoslavia to help provide reconnaissance and intelligence support for this mission was required on a 24 hour, seven day a week basis. At least five General Atomics Predator UAVs were deployed in mid- July 1995 to Albania and Croatia to support UN peacekeepers in some of the more hazardous reconnaissance missions.

The key arguments for expanding the use of UAVs is that they are cheaper than satellites, are far more flexible, can be operated below cloud cover when some satellites may be blind to the ground picture, and can often provide the imagery detail not available from satellites. Weather at this stage still restricts the operation of most UAVs and, in certain circumstances, they may need to rely on satellites or other relay vehicles for the transfer of data. Most UAVs use UHF or microwave links to carry real-time video signals. When sensor data cannot be transmitted by line-of-sight communications direct to the GCS, it must be stored on board or relayed to another platform (UAV, satellite, etc.) via another data link. These data links are regarded as a major weak point in UAV's due to their susceptibility to jamming and range limitations. If the data is not transmitted in real-time, the UAV needs an on-board data storage system from which the information can be retrieved either through later transmission to a ground station, giving a near-real-time transmission capability, or on the ground at the completion of the mission. Data storage systems can be substantial, depending on the sensor, amount of data collected and its ease of compression. For example, a high quality, high definition photograph collected digitally through an electro-optical system may require up to one gigabyte of storage.¹³ Data compression techniques enable mission data to be stored for a given storage capacity, however, they usually causes some loss of definition in the image.

The data storage requirements associated with UAVs will grow as their use becomes commonplace, and the endurance of these platforms increases. Over 60 UAV systems are presently available or are under full-scale development for the global market, with most systems

being the smaller, simpler vehicles employed at the tactical level.¹⁴ The latest Tier II-Plus and Tier III Minus UAVs being developed in the US for operational deployment in the post-2000 time frame, however, have the size and complexity of some manned aircraft but have far greater endurance. For example, the Tier II Plus Global Hawk high-altitude UAV scheduled to fly in late 1996 has a wingspan of over 35 meters, a length of 13.5 meters and a weight of 10,500 kilograms. A vehicle with Global Hawk's capability provides a surveillance coverage of 137,196 square kilometers or 1,900 spot targets per 24 hour period.¹¹ Tier III systems will be similar in size to Tier II Plus but will be stealthy and are planned to remain aloft for up to three months. Another capable but more survivable UAV system, developed by Lockheed Martin Company's 'Skunk Works', is the Tier III Minus DarkStar. But with both Global Hawk and DarkStar each having an estimated cost of US\$10 million, they cannot be considered readily expendable. This creates a need for significant on board storage for a detailed navigational database to enable the UAV to safely egress in the event of data link loss or other damage. With most of the emphasis for current UAVs being on the reconnaissance, surveillance and targeting acquisition (RSTA) roles, most of the payload developmental work has been on imagery sensors and the means to communicate this in real or near-real-time. As precision guided munitions require precision reconnaissance, much of the development has focussed on producing lightweight sensors that combine wide-area coverage with the necessary degree of resolution. For example, the TRA Global Hawk will have to carry an 820 kilogram package of EO, IR and SAR sensors with resolutions of 0.3 meter in the spot mode and 1 metre in the search mode, covering over 400,000 square kilometers per hour. Sensors cued with a GPS reference will be able to locate targets to within 20 meters.¹⁵

With the amount of data which can be collected by on-board sensors, there is a need when operating out of line-of-sight to transmit all or part of the data via satellites to commanders in real-time. It is now possible to process raw SAR data on board a vehicle rather than on the ground which reduces the necessary data link capacities to between one eighth to one tenth. The challenge for the future is to be able to transmit all of the data from all sensors in real-time or in periodic bursts. With the smaller Predator UAV, a program has already been tested using a direct broadcast transmitter to send 20 channels of digital video routed around the battlefield to allow commanders to 'channel surf' for imagery.¹⁶

The technological improvements and miniaturization of payloads is greatly increasing UAV capabilities to provide many more planning options. In the future, a UAV could be launched in one country, climb to high altitude and fly undetected over another country. Once in position, it could remain in a lazy orbit, unchallenged for periods of three months or longer. During this time it could observe, record, identify, track and fix any number of potential targets. Alternatively, it could intercept communications and electronic transmissions, relaying this information using microburst transmissions¹⁷ to an overhead satellite or other UAVs for analysis at a headquarters located anywhere on the globe. If that same UAV carried precision weapons to be launched once hostilities broke out or as a preventative or pre-emptive strike, the problems facing the defender increase enormously.

The URAV uninhabited reconnaissance aerial vehicles (URAVs), and space options are attractive as replacements for AWACS and Joint STARS. Both the AWACS and the Joint STARS use much of their volume for crew and displays, and loiter time is restricted by fuel consumption and crew limits. These systems offer the possibility of monitoring the entire world continuously at reasonably high resolution. Consider that the actual information content from a 10 m system is one bit per pixel spatial and 100 bits spectral. Both SAR and visible images assume that the total information content is 100 bits/pixel over the entire world once per hour. If we observe one percent of the world, about 1000km x 1000 km, at a rate of once per second, the data rate is 1.3 TB/s. Surveillance of all of Iraq at a rate of once per minute would require 5 GB/s. Obviously, the collection and storage of this sort of data throughput will require massive storage and analysis capabilities in the years to come.

3.2.2. Training and simulation: synthetic environments

The scope of the DoD Modeling and Simulation Vision extends from high-fidelity engineering models to campaign-level simulations involving joint forces, embracing education, training and military operations; analysis; research and development; test and evaluation; production and logistics, as shown in Figure 2-1.¹⁸

Advanced M&S may integrate a mix of computer simulations, actual warfighting systems, and weapon system simulators. The entities may be distributed geographically and connected through a high-speed network to create a *Synthetic Environments*, or internettted simulations that represent activities at a high level of realism from simulations of theaters of war to factories and manufacturing processes, which allows immersion into the environment being simulated. The DoD projects that the uses of Modeling and Simulation will include:

- Training for the complete spectrum of military operations for all regions of the world and affected regions of space. Exercise and training feedback will be available in near-real-time, with after-action reporting systems and exercise reconstruction systems providing a robust analysis capability. These reporting and reconstruction systems may require vast data storage capacities.
- To evaluate readiness, assess warfighting situations, and assist in the development and evaluation of operational plans, doctrines, and tactics. Decision makers will be able to simulate and evaluate the consequences of alternative courses of action during deliberate and crisis action planning on short notice via easily accessible M&S data resource repositories and automated scenario generation. Like the reporting and reconstruction systems, the scenario generators will need similarly large data storage systems.
- To allow warfighters and military planners to rehearse missions by immersing the warfighters in a synthetic environment that accurately simulates the anticipated terrain, environmental conditions, and threat.

Large complex warfighting simulations are already a reality,¹⁹ such as the 1997 Advanced Technology Demonstration called Simulated Theater of War '97 (STOW-97). The goals were to show the capabilities of distributed information system technology in a combat development and training exercise that will involve up to 50,000 simulated and non-simulated entities. DoD plans to have a working holodeck²⁰ by the year 2020.²¹ Virtual reality generates artificial situations that can be explored, "touched," and modified featuring all of the human senses and allowing the user to become immersed in the scenario. United States Special Operations Forces will be able to use this technology to quickly and accurately assemble a computer mock-up of any area in the world where they may have to make a forced entry or rescue mission.²² The Air Force will likely train pilots in a 3-D virtual reality simulator, complete with audio that mimics their current surroundings.²³ However, real-time information fusion is impossible without several orders of magnitude improvement in data storage capacity and processing speed.²⁴ Exponential increases in both these domains of computing power continue in the commercial sector. Rome Laboratory's C³I Technology Area Plan articulates the issues of storage and speed as a thread in virtually all of its thrusts.²⁵

Additional M&S Dimensions

- Level of Resolution
- Degree of Human Participation
- Degree of Physical Realism
- Time Management Method
- Time Step Resolution
- Degree of Distribution
- Computational Complexity

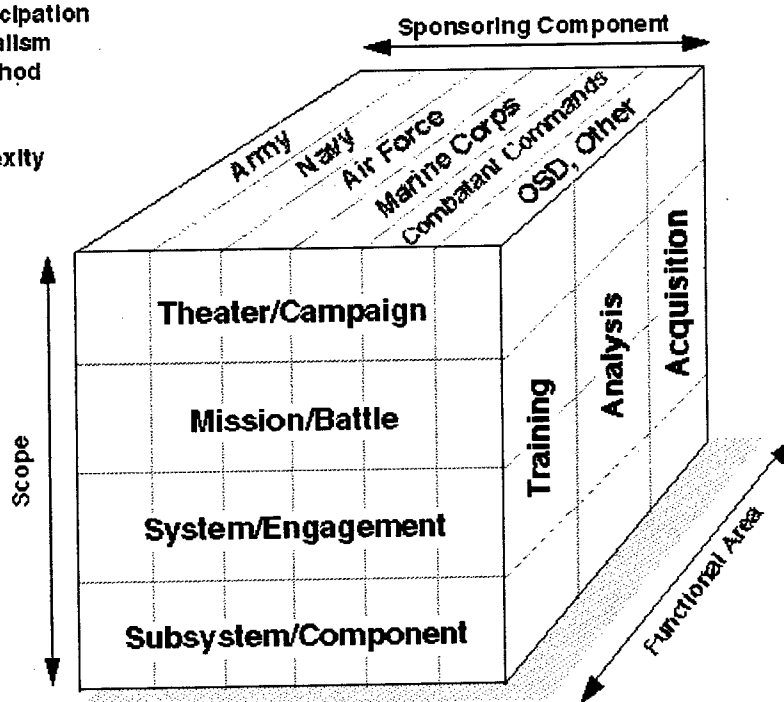


Figure 5. Range of M&S Embraced by the DoD M&S Vision
(from <http://www.dmsi.mil/docslib/mspolicy/msmp/1095msmp/Chap2.html#2A>)

Training enhancements made possible through technological strides in VR will have broad applications, from combat-environment simulation to emergency room training; the limits appear bounded only by lack of imaginative application.²⁶ A key feature of future training systems is the capability of accurately measuring the level of readiness for the individual, unit, or force as a whole. This requires fully automatic data collection and analysis for real-time,




Figure 6. Goddard Earth Observing System (EOS) Data Assimilation System 1998 hardware requirements and projections. (from reference 30)

continuous, and adaptive assessment of performance in various scenarios. Again, the continuous logging and analysis of these training simulations will necessitate a large on-line data store.

3.2.3. Mission support: Transportable Training and "World Map"

Computer-based multimedia training (CBT) is rapidly evolving due to its promise of higher efficiency training at reduced costs, and the growing availability of most of the required hardware. Examples of CBT include prerecorded lessons interactively served to students across

a widely distributed computer network, interactive lessons or manuals distributed on disk to enable remote use of the information, and recorded mission training and rehearsal sessions that may then be carried along to support real-time decisions during the mission. Training for the Air Force's new generation Advanced Technology Fighter, the F-22, has a large CBT component, supporting animation, graphics, and digital photography and video, requiring 400 GB of storage, and up to 3 Gb/s data transfer rates.²⁷ Similar systems are being considered for training AWACS personnel. Additionally, since many of these platforms are being, or will be used by international allies, the requirements for multiple language editions and ease of distribution around the world further the need for high capacity, high performance optical storage.




Figure 7. Potential military applications for optical storage systems: projected data capacity requirements.

Mission planning and analysis is another storage-hungry application. The Airborne Tactical Information Management System (ATIMS) program has been demonstrated on Apache, F-14, and B-1B aircraft with the goal of reducing pilot workload and systems avionics costs, and enhancing mission flexibility, effectiveness, and survivability.²⁸ It will enable rapid preplanning and integrated rehearsal and reporting by allowing the pilot to plan and rehearse a mission on the ground (presumably with multimedia and simulator support). The pilot then carries a storage device containing that planning session on the plane during the mission to provide situation display aids and enable in-flight re-planning, mission management, and embedded training. Additionally, the storage device can record information and signals received

by the pilot during the mission, to be available for mission debriefing, as well as future training and planning.




Figure 9. Potential military applications for optical storage systems: projected data access time requirements.




Figure 8. Potential military applications for optical storage systems: projected data rate requirements.

In the next century, one Air Force goal is for all aircraft to have the navigation database to fly anywhere, anytime, on any route independent of external data.²⁹ Precision mapping should be to equip each aircraft and planning system with a map of the entire world to one-meter accuracy, which will require 10-20 terabytes of data storage with suitable compression. The "onboard world" will enable the ultimate in moving map navigation and self contained undetectable terrain avoidance. The Defense Mapping Agency currently provides products to support such an "onboard world", including Digital Terrain Elevation Data (DTED) and Digital Feature Analysis Data (DFAD) which can enable tree-top level flying by providing terrain information, including man-made structures such as TV towers and high-tension wires. Increased storage capacity would enable the information to be displayed as high-resolution multispectral imagery. This information would enable combat aircraft to fly directly to the target at low altitudes using only DMA data and GPS navigation aids, in all weather, without terrain following/avoidance radars that might alert the enemy to their presence. These capabilities are particularly important for the F-15, F-16, F-111, B-1, and B-52 aircraft.

In 1991, NASA launched a comprehensive program to study the Earth as an environmental system, called the Earth Science Enterprise. Satellites and other tools are used to intensively study the Earth, to gain a better understanding of the global environment and how the Earth's systems of air, land, water, and life interact with each other. Such studies will yield improved weather forecasts, tools for managing agriculture and forests, information for fishermen and local planners, and, eventually, the ability to predict how the climate will change in the future. The collection and analysis of the data from this Earth Observing System (EOS)³⁰ is one of the Grand Challenges of computing. Special observational programs are used to study problems in weather forecasting, atmospheric chemistry, the influence of the tropical oceans on global weather, and ozone depletion. EOS data is assimilated at at least 10 distributed active archive centers (DAACs), and the computing and mass storage requirements were found for the Goddard Earth Observing System (EOS) Data Assimilation System (GEOS DAS). As shown in Figure 6, typical input storage requirements are 2 TB of on-line disk storage and 8 TB of near on-line storage (currently tape silo). Short term research output as well as archived research output will require a total of 50 TB by 1998 and 200 TB by 2002.

There is a large spectrum of both military and commercial applications that require data capacities exceeding 100GB per platter while supporting data rates in excess of 1Gb/s, as shown in Figure 7 - Figure 9. In addition to the application examples discussed above, there are considerable data storage requirements for continuous logging of instrumentation data streams and airborne video imagery recording (of Head-Up-Display and cockpit video) in the Joint Strike Fighter (JSF) and B-1B.

3.3. Potential synergies between military and commercial roadmaps




Figure 10. Potential overlap of mass-market and military applications.

As supercomputing power and very large databases become more common in the commercial sector, the data storage requirements for military and commercial applications are becoming increasingly similar, and the potential exists for significant cost saving by using commercial storage products in military environments. One potential difference in application requirements is the need for ultra high reliability in harsh environments. However, these large databases are now becoming the lifeblood of many businesses, and this fact, together with the increasing use of portable electronics is accelerating the introduction of techniques to dramatically improve robustness and reliability in data storage hardware. Improvements in VLSI electronics and increasing storage capacity has enabled sophisticated error control algorithms in the drives, using electromagnetic tests and digital coding techniques to detect and correct data reading errors. Today, these algorithms are being extended to predict drive reliability problems in advance of drive failure (beyond the capacity of the error control algorithms). This predictive reliability technology, termed "SMART" ("Self Monitoring and Reporting Technology") is now an industry standard for monitoring drive performance and reporting incipient failures, so that the drive's user can take action before actual failure occurs.³¹ This reliability issue is becoming critical for magnetic hard drives. Magnetic hard drive technology is being pushed to its limits, and most computer manufacturers are skeptical of the claimed million hour mean-time-between-failure.³² This presents a great opportunity for optical storage with its inherently greater reliability.

As discussed above, data mining and data recording/archiving applications in both the military and commercial domains often share similar requirements. Additionally, emerging commercial applications in collaborative telepresence (see Figure 11), virtual reality gaming, and commercial computer based training are likely to develop requirements along the lines of military battle simulation and training applications. U.S. companies will spend \$60 billion this year on training their employees, and increasingly that process is being conducted over the Web, according to International Data Corp. Technology-based training systems have grown 40% annually in recent years, and analysts expect that trend to continue. "This market is really taking off, and tremendous attention is being paid by venture capitalists and other private investors,"

says the editor of an online training newsletter. One producer of online training software points out that offering online training in such sensitive subjects as sexual harassment provides a level of privacy not found in classroom training programs. "The privacy of the computer makes people more psychologically open. People would rather go to the computer than talk about it in a group training session."³³

USC
UNIVERSITY
OF SOUTHERN
CALIFORNIA



Figure 11. Collaborative telepresence system concept.

Another synergy between military and commercial optical data storage applications can be seen in the effect that optical storage may have on driving a nation's optoelectronic industry. As we enter the "optical information age", optoelectronics will be an increasingly important foundation of larger industries, as shown in Figure 12.




Figure 12. 1995 economic impact of optoelectronics (adapted from OIDA).³⁴

The Optoelectronics Industry Development Association (OIDA) forecasts explosive growth in the optical storage market from \$6 billion to \$50 billion by the end of the next decade (2010),

mostly supported by video- and computing-related products. Currently Japanese optical storage makers dominate the world market and have been successful in building an enormous technology base in optoelectronics to take advantage of this explosion. Indeed, optical storage manufacturing has contributed significantly to the establishment of an optoelectronics infrastructure in Japan, an infrastructure now available for further development of optoelectronic components in general and optical data storage in particular.

4. Conventional Storage Technology

4.1. Optical disk hardware review

Due to their cost advantage at high capacities, robustness, lifetime, and removability, optical storage systems are commonly used for very large storage systems and backup systems ("jukeboxes"). The advantage of optical systems for this market is that they have much shorter access times than tapes and are much more wear free. Optical storage, due to its immunity to head crashes and media wear, is often preferred to magnetic storage in many military applications; for example to transfer data from the tactical mission planner to a combat aircraft.

The storage media of most optical storage systems in production are in the form of a rotating disk. Figure 13 and Figure 14 illustrate the structure and functionality of a typical optical disk system. In general, the disks are preformatted using grooves and lands (tracks). This enables the positioning of an optical pickup and recording head to access information on the disk. A focused laser beam emanating from the optical head records information on the media by changing the material characteristics. To read a recorded bit, the laser generates a small spot on the media. The presence of a written bits modulates the phase, intensity, polarization, or reflected power of the readout optical beam which is subsequently detected by a photodetector in the optical head. Drive motors and servo systems rotate and position the disk media and the pickup head, thus controlling the position of the head with respect to data tracks on the disk. Additional peripheral electronics are used for control and for data acquisition, encoding, and decoding.

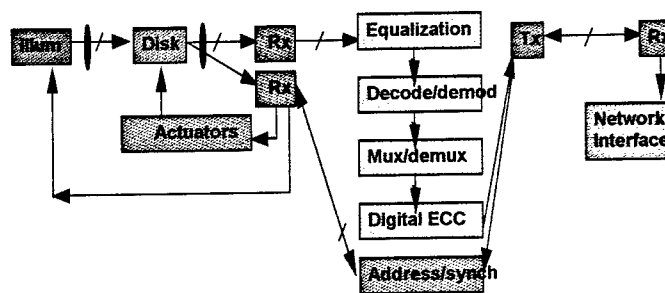



Figure 13. Conventional optical disk pickup head architecture. Figure 14. Functional diagram of an optical head.

As for all data storage systems, their storage capacity, data transfer rate, access time, and cost characterize optical disk systems. The storage capacity of an optical storage system is a direct function of spot size (minimum dimensions of a stored bit) and the geometrical dimensions of the media. A good metric to measure the efficiency in using the storage area is the areal density (Gb/sq. in.). Areal density is governed by the resolution of the media and by the numerical aperture of the optics and the wavelength of the laser in the optical head used for recording and

readout. Areal density can be limited by how well the head can be positioned over the tracks; the track density (tracks/in.) measures this. In addition, areal density can be limited by how closely the optical transitions can be spaced; the linear density (bits/in.) measures this. The data transfer rate of an optical storage system is a critical parameter in applications where long data streams must be stored or retrieved, such as for image storage or backup. Data transfer rate is a combination of the linear density and the rotational speed of the drive. It is mostly governed by the optical power available, the speed of the pickup head servo controllers, and the tolerance of the media to high centrifugal forces during high-speed rotation. The access time of an optical storage system is a critical parameter in computing applications such as transaction processing; it represents how fast a data location can be accessed on the disk. It is mostly governed by the speed of the head movements and is proportional to the weight of the pickup head and the rotation speed of the disk. The cost of an optical storage system is a parameter that can be subdivided into the drive cost and the media cost. Cost strongly depends on the number of units produced, the automation techniques used during assembly, and component yields.

As a starting point for analysis of future optical storage progress, the performance capabilities of conventional data storage systems at all levels of the storage hierarchy were surveyed. This data is shown in Figure 15, which highlight three main "gaps" in the performance hierarchy. First, no current technologies simultaneously support high capacities and low access times. Second, high data transfer rates are unavailable from single drives, and require the use of semiconductor memories. Third, the progress of peripheral connections to the storage units is outstripping the growth of the transfer rates of the units themselves.




Figure 15. Performance comparisons of conventional data storage technologies (1997).

Optical storage R&D typically concentrates on the following efforts: reducing spot size using lower-wavelength light sources; reducing the size and weight of optical pickup heads using microoptics, diffractive optical elements, and advanced micro-packaging techniques; increasing rotation speeds using larger optical power lasers; improving the efficiency of error correction

codes; and increasing the speed of the servo systems. Equally active R&D efforts, especially in Japan, are focused on developing new manufacturing techniques to minimize component and assembly costs. Japanese manufacturers are convinced that over the next decade new emerging applications will pull the performance of optical storage systems and that evolving conventional optical storage technologies (DVD and MO) are capable of satisfying these demands for at least another seven to eight years.

4.1.1. Phase change media projections

A potential roadmap for DVD systems using phase change media is shown in Figure 16. During 1998, phase change (PC) storage system manufacturers were heavily involved in the advanced R&D of 4.7GB DVD-RAM products. PC manufacturers may combine higher NA optics with shorter wavelength blue lasers and Single Carrier Independent Pit Edge Recording (developed by SONY) together with Radial Direction Partial Response (RPR) encoding technique to achieve the 15GB capacity double layer disks required for the HDTV standard. Experiments are also conducted as well with more transparent PC media layers to enable four-layer disks to be used in higher capacity products. However, even with the incorporation of these techniques the effective areal density (including multiple layers) of phase change media is expected not to exceed 50Gb/in² because of diffraction of light (limiting the areal density per layer) and optical crosstalk (limiting the number of layers). To overcome these limits significant developments must occur such as the development of very low cost miniature lasers operating with a wavelength of ~250nm. At this point in time these limits appear very hard to overcome, thus windows of opportunity for alternative low cost data storage techniques are expected to appear.




Figure 16. Example roadmap describing performance potentials of phase change media
(Source: Hitachi (HCRL))

4.1.2. Magneto-optical media projections

A similar situation exists for MO storage. As an example, consider the roadmap provided by Fujitsu as shown in Figure 17. MO drive manufacturers now believe that the MO technology can be extended to 5.2 GB capacity in the near future, and then to 10.4 GB on a 5.25" double-sided disk using 680 nm wavelength and a lens with an NA of 0.55. Beyond that, they plan to rely on one of the magnetic super-resolution (MSR) techniques, either Magnetic Amplifying Magneto Optical Systems (MAMMOS) from Hitachi-Maxell or Domain Wall Displacement Detection (DWDD) from Canon, in addition to magnetic field modulation, to enable further capacity increases. By 2002 they may be able to achieve areal densities approaching 20Gb/in² by combining one of the MSR techniques with the use of a blue laser and larger NA optics. By adopting a format similar to DVD, MO researchers at Fujitsu are contemplating 36GB capacity disc systems as VCR replacements. Finally, they envision using SIL lens and parallel heads to extend the areal density and data rate of MO products to approach 100Gb/in² and 1Gb/s respectively. However, reaching areal densities beyond 100Gb/in² is becoming increasingly difficult and unlikely unless compact ultra short wavelength lasers become available.




Figure 17. Example roadmap describing performance potentials of MO media based systems
(Source: Fujitsu)

4.2. Channel Electronics and Peripheral Interconnection Trends

Many emerging applications need "firehose" access to massive data storage, and the microprocessor system architectures and I/O interfaces to provide these high bandwidth connections are becoming available. For example the 32 bit wide, 33 MHz PCI bus already offers 1056 Mbit/s data rates, switched bus architectures such as the Mercury RACEway or Sun Microsystems's S-Connect can provide highly scalable system-area-network transfer rates of 1280 Mbit/s and 1760 Mbit/s. Multi-fiber networking technologies like PONI are emerging to offer transfer rates of 6240 Mbit/s over LAN and enterprise server distances. Given the 11 Mbit/s transfer rates of DVD systems (typical to most high capacity storage solutions), the need to bridge this widening gap is clear. Even the latest expensive RAID 0 systems configured for

maximum data rate provide only about 800 Mb/s, still falling below current PCI bus speeds. This section considers this gap by examining the trends in internal data transfer rates and peripheral interconnection rates of disk-based storage systems.

4.2.1. Channel Electronics

In the past few years, the *internal data rate* in rotational disk based storage devices such as magnetic hard disks (HD), digital versatile disks (DVD) and magneto-optical (MO) disks has increased at a rate of 40% per year. These devices utilize a serial read channel for data detection and error correction whose performance is linked to the underlying implementation technology (currently CMOS or BiCMOS) and the channel architecture. SIA predicts that CMOS gate frequencies will continue on a 9% annual growth rate over the next ten years. To overcome the gap in performance growth between CMOS and internal data rates, architectural enhancements will be required, such as more complex channel detection circuitry, more powerful error correction coding (ECC) and parallel read-out of data channels. This has certainly been true for magnetic HD in recent years, where a shift from peak to PRML detection has resulted in a 33% increase in overall performance. This section describes the trends in secondary storage device data rates and read channel performance, focusing on disk based technologies, both removable and non-removable.




Figure 18: Read channel functional blocks for disk based storage devices.

Figure 18 shows the read channel functional blocks for a magnetic HD. For optical storage, the disk interface is an optical detection module. The disk interface passes an analog signal detected from the disk to the read channel. The read channel contains four function blocks. The first performs automatic gain control and low pass filtering to reduce channel noise and is typically implemented with analog circuitry. The equalization is either performed in analog or digital format, depending on the implementation. The digital implementation includes an ADC with FIR equalization. Similarly, the data detection can be performed with an analog detector, as in the peak-detection method, or digitally, as in more advanced sequence detectors. The decoder block typically implements RLL decoding, which combines groups of bits into a symbol for error correction using Reed-Solomon ECC. The RLL and ECC encoding both account for a reduced coding efficiency (e.g., ratio of user bits/detected bits). Magnetic HD typically implement RLL with a code rate of 8/9 and a 97% efficient RS code for an overall efficiency of 86%.

Removable optical disks have a much lower code efficiency, for example the DVD 56% of the data read is overhead.

Recently, magnetic HD has shifted from peak-detection to partial-response maximum likelihood (PRML) data detection. This shift will occur also in MO and DVD channels. The PRML detector compares a series of data samples taken at regular intervals to all possible bit patterns to find the most likely series of bits (code symbol) encoded on the disk. Bit sequence detection differs from peak-detection, which detects each bit one at a time without information from previous bit values. Current PRML methods allows for a lower (1.5-3-dB) signal-to-noise ratio (SNR) for the same output BER and allows for a 30% linear-bit-density increase on disk surface. PRML's higher SNR/BER performance also allows for an increase the RLL code rate. Peak-detection operated at a 2/3 code rate, where today's PRML channel use a code rate of 8/9. The digital circuitry scales better than analog circuitry in terms of silicon area with advances in IC technology.




Figure 19: Comparison of read channel transfer rates with peripheral interconnect (PI) technology. Commodity PI follows the trend in desktop interconnects, such as SCSI and Firewire. High performance PI follows the HIPPI standards for interconnecting workstation/supercomputer clusters.

4.2.1.1. Future Trends

The future performance of serial read channels will depend on IC technological enhancements and architectural developments. Figure 19 shows the trends for serial read channel bandwidth, commodity level peripheral interconnects (following SCSI and Firewire predictions) and high performance peripheral interconnects (following HIPPI predictions). The read channel curve was derived by following the previous three generations (implementing PRML) and continuing the trend to 2010. The finding is that the growth rate follows CMOS gate frequencies (9% annual growth). This curve assumes that future read channel complexity will be offset by future architectural/layout enhancements. Thus the underlying CMOS technology will provide performance growth. The read channel complexity is described in the bottlenecks section below.

These plots show an increasing gap between the interconnect and serial read channel performance.

Figure 20 shows compares the internal data rates for magnetic HD, RAID, DVD, MO and near




Figure 20: Projected internal data rates for various storage technologies. Internal data rate is derived by multiplying the linear bit density by the rotational velocity of the disk.

field magnetic technologies. These technologies all share a similar type of serial read channel. Future growth will be due to increasing the linear bit density stored on the disk and increasing the disk RPM. Magnetic HDs are scheduled to reach the superparamagnetic limit near the year 2005. Another recording method, such as perpendicular or contact recording would be required for further advancements. DVD and MO will advance as shorter wavelength lasers become available, Toshiba is currently demonstrating blue-purple laser at 417 nm. Eventually the NA requirements will be difficult to maintain. Terastor's near field recording technology expects no such limitations on density for the next ten years. The RAID curve plots a four HD raid system. Beyond four HDs in RAID systems the transfer rate saturates due to the large latency of magnetic HDs. However, future RAID configurations (called round robin) will allow a tradeoff between transfer rate and overall latency.

4.2.1.2. Bottlenecks

The speed-limiting building blocks for digital read channel implementations are the FIR equalization filter, the analog-to-digital converter (ADC) and the PRML sequence detector. The sequence detection block is the most complex design and ultimately will determine future scaling of serial channel data rates. The PRML detector typically uses a Viterbi detector to perform recursive calculations using current and previously stored data samples. This recycling of past results creates the speed bottleneck in the add-compare-select unit (ACS) within the Viterbi detector. There are two types of PRML detectors currently in the market. The *partial response-class-IV* (PR4) detector uses two ACS units and the *extended PR4* (EPR4) uses eight ACS units. The EPR4 offers a better signal-to-noise ratio and thus allows higher bit densities.

Year	CMOS Feature (micron)	Gate Delay (ps)	Channel Data Rate (Mbs)
1989	1.2	380	130*
1992	0.8	290	180*
1995	0.5	210	250*
1998	0.25	120	550
2001	0.18	92	660
2004	0.13	70	880
2007	0.10	54	1100
2010	0.07	41	1210

Table 2: Channel electronics roadmap for serial read out. The * indicates commercially available implementations by Silicon Systems, Inc. and Marvell Semiconductor.

The ultimate choice of sequence detection method depends on the ECC capabilities required and on the pulse shape characteristics of a given hard drive, which is determined by the head-media performance and linear and track density. Detectors can be classified by the lowest SNR tolerated in obtaining a given BER at a given user-bit-density (D_u) (defined as the pulse width due to magnetic transitions divided by the bit clock period). As the achievable user-bit-density progresses from the current 2-2.5 range to 2.5-3, demand for more SNR efficient solutions may introduce shift from PRML detectors to decision feedback equalizer (DFE) and fixed-delay tree search (FDTS) architectures³⁵. For example, at a BER of 10^{-6} , the PR4 offers lowest SNR requirements up to D_u of 1.4. For a D_u ranging from 1.6 to 2.2 the EPR4 is the best, and for D_u greater than 2.4 an ML detector based on FDTS is superior. The DFE and FDTS have a speed advantage, since these architectures reuse past decisions only (rather than a number of partial results).

4.2.1.3. Parallel Read Channels

If disk storage densities and peripheral interconnect transfer rates continue increasing at their present rate, the read channel will likely become the limiting bandwidth factor for storage I/O. Current convolutional detection methods and ECC will become more complex as the increased linear density will decrease SNRs and require more extensive error correction. This added complexity will require more performance from the CMOS. Several methods exist for achieving 'one-time' gains, such as changing the fabrication technology or finding a more optimal convolutional algorithm. However, adding parallelism to the read channel is a scalable solution. One method of introducing parallelism to the read channel is to use a time multiplexing method combined with buffering. Since current and future detection methods perform data sampling before performing the data detection (where the ACS bottleneck occurs), samples could be time multiplexed between two or more read channels operating at a lower clock rate. If servo methods allow, another parallel method is to provide each read head with a dedicated read channel. This possibility exists in magnetic HD which contain several independent heads. In optical storage methods, several readout beams could be supplied to implement parallel readout. Several issues must be addressed for parallelism, such as advanced servo mechanisms for controlling multiple read heads, recovery of multiple clocks, parallel implementation of convolutional and ECC algorithms.

4.2.1.4. Summary

If trends established over the last few years continue the data rates for read channel electronics are expected to grow at rate roughly equal to CMOS technology performance, (roughly 9% increase in throughput per year). This trend currently applies to all disk based storage as linear densities and spin rates increase. Table 1 shows the progression of IC technology as given in the




Figure 21. Channel electronics roadmap for serial read out.

SIA roadmap and follows current read channel implementations to the year 2010.³⁶ Over the last

3 years improvements in channel electronics have accounted for major gains in the linear density of usable data on the disk surface. The shift from analog peak-detection schemes to digital PR4 sequence detection accounted for the ability to detect a 30% increase in linear-bit-density and a 33% increase in code efficiency. In the next few years, the code rate shift from 8/9 to 16/17 (0.94), is expected and further gains gained by increasing the code rate efficiency will be minimal. As linear-bit-density increases, more sophisticated sequence detection will be required. Efforts are underway to increase the complexity of the maximum-likelihood detectors (increasing the number of ACS units), using more advanced DFE or FDTS architectures and/or by increasing the coded symbol length. These advanced detector architectures will allow detection of higher linear bit densities and thereby increase the read-out rates. However, in order to gain throughput, advances in the fabrication technology must offset the increased complexity of the detector.

4.2.2. Interconnection Trends

The anticipated service demands of the next century place strict constraints on future networks and storage devices. Multimedia services such as Video-On-Demand (VOD), high definition television (HDTV), and video conferencing demand high data rates and capacities, short access times, and a guaranteed Quality of Service (QOS). To meet these demands, many high-speed network platforms have been developed and high speed, large capacity storage technologies have been and continue to be researched extensively. As these technologies advance, however, a bottleneck is being created at the interface between the network and the storage device.³⁷




Figure 22. Network-to-storage system bottleneck: performance requirements for various network-attached services (from reference 37).

This section describes performance trends in peripheral interconnects during the past 10 years and projects these trends over the next 10 years. The interconnect technologies discussed here are applicable to storage devices either now or perhaps in the future. These interconnects may serve as links between processor and storage devices or within storage clusters.

Peripheral interconnects (links less than 300 meters, i.e., desktop to LAN) can be divided into two classes: *high performance* link protocols that define switched networks with very low latency and *cost/performance* (commodity) link protocols defined on loop or daisy-chain networks which trade latency for expense. Advances in IC technology and computing architectures will increase the aggregate data I/O of peripheral devices by increasing both channel bandwidth and number of channels (e.g., bus width) per device. The channel bandwidth limitation will not be the physical link capability (most protocols extend to fiber links whose capacity, even on a single fiber, is much greater than will be demanded to the year 2010), but rather the cost of the electronics implementing the protocols and performing serializing/deserializing functions. These costs will determine the trend in parallelism in physical links as channel rates surpass the 100 Mb/s mark.

HighPerformanceLinks - High performance links are typically routed on switching fabrics to reduce latency. The Semiconductor Industry Association report predicts low latency ($< 1\mu\text{s}$) and very high aggregate bandwidth ($> 300\text{ Gb/s}$) requirements by the year 2010. For high performance links, the switching fabric implementation determines overall bandwidth and latency. The main issues concerning parallelism are the time required to serialize/deserialize channels and channel skew. The state-of-the-art for parallel electronic channel is the micro-coax (the Gore Electronics cable used in SCI links) which exhibits 50 ps/m of skew. However, fiber ribbons have demonstrated a channel skew less than 1 ps/m.

Cost/Performance Links - Serial fiber links are capable of supplying more bandwidth than that required for most commodity applications such as workstation station/pc clusters and distributed memory access systems. Latency is not critical as long as it is guaranteed (and thus accounted for on the receiving end), as in isosynchronous services (e.g., digital video). Serial link will remain popular (FireWire, FC-AL) because of the smaller lighter cabling. However, eventually the high cost of multiplexing/demultiplexing wider, faster busses into a serial channel will force some parallelism.

4.2.2.1. Survey of Peripheral Interconnects

This section surveys the trends in peripheral interconnect (PI) technologies over the past few years and predicts performance growth for the future.

4.2.2.1.1. High Performance Links

High Performance Parallel Interface (HIPPI)

HIPPI is an ANSI standard that defines physical layer and protocol for switched network interconnect scheme. HIPPI is best suited for supercomputer connectivity, high-end workstation and storage device clustering, and as a backbone connecting other networks. HIPPI was created

in the late 1980s, when Los Alamos National Laboratory in New Mexico needed a connection-switched interconnect that would run at 800 Mbps between its Cray Research Inc. supercomputers and visualization devices. This original standard, HIPPI-800, creates 32 parallel data channels between nodes on 50 shielded twisted pair copper lines, with each channel operating at 25 Mbps. Switches up to 32 x 32 are commercially available. In 1991, HIPPI-1600 (dubbed 'Wide HIPPI') standard increased the number of parallel data channels to 64 to achieve 1600 Mbps throughput between nodes. Both HIPPI-800 and HIPPI-1600 are optimized for streaming of large volumes of data, not for low-latency dynamic interactive usage.

An ANSI working group formed in 1996 formed the HIPPI-6400 standard (also called 'SuperHIPPI'). HIPPI-6400 provide 6.4 Gbps duplexed transmission using 32 parallel data channels over copper up to 50 meters or fiber to a kilometer. HIPPI-6400 has many improvements over previous standards, including shorter packet lengths (the cell length is 32 bytes), error correction, and hardware flow control. These improvements reduce the latency to less than 1 microsecond. Silicon Graphics Inc. is currently testing the first HIPPI ASIC, SuperHIPPI Media Access Controller or SUMAC, which will be the basis of HIPPI products. Example HIPPI-6400 products due this year are Genroco International's five port hub for PCI, Sbus and FibreChannel networks; and Essential Communication's 32 x 32 non-blocking crossbar. It is expected that future HIPPI products will provide throughput up to 12.8 Gbps by 1999 and 25.6 Gbps by 2000.

Fibre Channel (FC)

Fibre Channel defines a set of ANSI standards that define the physical media and transmission rates, encoding scheme, framing protocol and flow control, common services, and the upper-level protocol interfaces. Fibre Channel (FC) is designed for serial interconnects between workstations, mainframes, supercomputers, desktop computers, storage devices, displays and other peripherals. It is optimized for streaming of large volumes of data in storage area networks. In fact the standard include disk control features for high speed disk I/O. Supports mapping of several other protocols onto the Fibre Channel links, including HIPPI, ATM, SCSI and TCP/IP. The current standard (released in 1994) specifies rates up to 1 Gbps over distances of 24 meters in copper or 10 km on fiber. ANSI working groups are currently working on extensions to 4 and 8 Gbps.

There exist three types of network implementations of FC, a point-to-point interconnect defined by FC-GF (generic fabric), a switched network FC-SW, and a loop architecture FC-AL (arbitrated loop). FC-SW and FC-GF are higher performance, lower latency implementations for higher-end applications that can link up to 16 million nodes. FC-AL is a token loop network (although usually implemented with a single hub), which allocates bus bandwidth across nodes (up to 126 nodes can be connected). FC-AL was developed as lower cost implementation of FC to compete with SCSI, specializing in disk I/O applications. Compared to SCSI FC-AL offers higher performance, greater number of nodes and longer link length.

Scalable Coherent Interface (SCI)

SCI was designed for high-performance highly parallel multiprocessors, but scales down to uniprocessors as well in order to get the economic benefits of high volume production. Contains protocols to maintain cache coherence in distributed shared memory systems. Uses shared memory model (i.e., a single address space for specifying data) rather than packetizing data to reduce latency. This protocol is optimized for random access transfer of small data packets in networks requiring guaranteed delivery. The current ANSI standard published in 1993 specifies 1 Gbps links, extensions to 8 Gbs are under study. Typical performance is currently in the range of 200 MByte/s/processor (CMOS) to 1000 MByte/s/processor (BiCMOS) over distances of tens of meters for electrical cables and kilometers for serial fibers. SCI/LAMP was designed to be interfaceable to common buses such as PCI, VME, Futurebus, Fastbus, etc., and to I/O connections such as ATM or Fibre Channel. The IEEE standard (P1596.8) specifies the connector and copper cable assembly for a 16 bit wide SCI link operating at 1 Gbps.

Asynchronous Transfer Mode (ATM)

ATM development originally targeted WANs for telecommunications applications and is currently being developed for large enterprise networks, LANs and backbone links between desktop networks. ATM LAN Emulation (LANE) standard supports multiple protocols, including Fast Ethernet, Token Ring, and FDDI. Latency was not an issue for WAN applications and telecom switches typically have greater than 10 microsecond latency. ATM is defined on a serial channel that can be implemented fiber such as on Synchronous Optical Networks (SONET) or on coax cable. In 1986 the standard for 155 Mbps ATM was released. The latest standard is 622 Mbps, released 1996. The ATM forum is currently working on extending the bandwidth to 2.4 Gbps and 10 Gbps. Future capabilities will include multimedia, quality of service support for both voice and video transmissions, and virtual networking.

Parallel Optical Link Organization (POLO)

POLO is a consortium of Hewlett-Packard, AMP, Du Pont, SDL and the University of Southern California developing high performance parallel optical data links for computer clusters, multimedia and switching systems. These parallel optical links may replace parallel electrical links found in SCI and HIPPI-6400 links described above. POLO demonstrated 10 parallel transmit and receive channels each operating at 1 Gbps for a 20 Gbps aggregate bidirectional throughput over 300 meter links. This effort is similar to parallel optical link technologies being developed by Motorola (OPTOBUS) and Optabahn and NGK. HP is expected to release a product based on POLO, called NPLEX, in late 1998. The first NPLEX links will be a 4 channel receiver or transmitter package with each channel operating at 1.25 Gbps. By mid 1999, HP expects to expand the number of parallel channels to 8 and 12.

4.2.2.1.2. Cost/Performance Links

Serial fiber links are capable of supplying more bandwidth than required for most commodity applications such as workstation station/pc clusters and distributed memory access systems. Latency is not critical as long as it is guaranteed (and thus accounted for on the receiving end), as in isosynchronous services (e.g., digital video). Serial link will remain popular (FireWire, FC-

AL) because of the smaller lighter cabling. However, eventually the high cost of multiplexing/demultiplexing wider, faster busses into a serial channel will force some parallelism.

SerialBus IEEE Standard 1394 (FireWire)

FireWire is targeted for consumer computer peripherals such as in multimedia computer environments, including video camcorders, digital VCR tape decks, and digital-video-disk (DVD) devices. Each link is a thin copper-based cable containing two twisted-pair signal lines, one for data and the other clock, and power lines. The copper links can be up to 4.5 meters in length. There are extensions to multi-mode fiber which can operate over 300 meter links. The network behaves like a single bus, multiple links are not independent. The latency is a maximum 125 microseconds and is guaranteed for isosynchronous service. The first standard defined 100, 200 and 400 Mbps links, the current standard is 800 Mbps and work on 1.6 Gbps and 8 Gbps links is in progress.

Serial Storage Architecture (SSA)

SSA is a derivative of an internal IBM serial SCSI protocol, that is being developed by a task group of the SCSI committee. SSA is optimized for links between computers and storage devices such as disk drives, CD-ROMs and tape drives. SSA links operate on 4 parallel channels, either shield twisted pair copper to 40 meters or fiber to 2.5 km. The 1995 version operated with 320 Mbps aggregate throughput, the current version operates at 640 Mbps.

Small Computers System Interface (SCSI)

The first SCSI standard of 1986 used 8 parallel channel to provide 32 Mbps throughput. Today, the Ultra-2 LVDS (low voltage differential signaling) standard uses 16 channels to provide 640 Mbps. Ultra-3, due in the year 2000, is targeting 1.6 Gbps throughput. Other than performance enhancements, work has focused on reducing the physical size of the cabling and connectors (a SCSI cable contains 68 conductors), reducing the complexity of usage, increasing the number of nodes (up to 60), and increasing the link length to 12 meters. SCSI remains targeted towards desktop peripheral applications.

Ethernet

Ethernet was developed in the late 70's operates over shared serial channel medium with a maximum throughput of 3 Mbps. Since the standard was published in 1985, the standard has grown to include new media systems for 10-Mbps Ethernet (e.g. twisted-pair media), as well as the latest set of specifications for 100-Mbps Fast Ethernet (1994). Current work the IEEE committee is working on standardizing 1 Gbps Ethernet.

4.2.2.2. Trends in Transfer Rates




Figure 23: Aggregate throughput for interconnect technologies applicable to storage devices. Data points represent commercial products. The graph highlights the trends in high performance versus commodity interconnects.

Figure 23 compares the aggregate transfer rates for both high performance and commodity PI technologies. This plot covers PI implementations from 1986 to current, as well as company predictions for future products. The data shows that high performance links exhibit roughly ten times the throughput of commodity links for any given year. For both groups, the throughput increases at a compounded annual growth rate of 32% (shown by the two lines in the graph). These growth rates match fairly closely with the SIA roadmap describing computer I/O rates, in which gains in transfer rate are due to a combination of increasing bus clock rates and increasing bus widths.³⁶ Continuing this trend to the year 2010 shows over 100 gigabytes/sec for high performance links and over 10 gigabytes/sec for commodity links. According to the OIDA roadmap, the aggregate bandwidth growth rate of high performance links can be maintained by parallel optical fiber links³⁸. This roadmap predicts data rates greater than 40 gigabits/sec per channel by the year 2010.




Figure 24: The number of data channels used for various PI interconnects versus the year of implementation.

The trend in parallelism has not been well established. This is due to the fact that channel bandwidth limitation is not the physical link capability (most protocols extend to fiber links), but rather the cost of the electronics implementing the protocols and performing serializing/deserializing functions. Figure 24 shows the trends in parallelism for various PI technologies. The decision to use parallelism lies mainly in the cost per link scaling over time. The cost of physical interfacing scales poorly over time. The cost of manufacturing connectors and cables show only a 2 to 4-fold cost reduction over the product lifetime. These links should be made as inexpensive as possible. In contrast silicon scales very well with experience, typically having a 10 to 100 fold cost reduction over 10 years. Therefore the general trend is to




Figure 25. Gap between peripheral interconnection rates and serial read channel rates (high performance is an extension of HIPPI, commodity is an extension of SCSI, serial read channel is an extension of PRML implementations).

push the complexity of the link into the IC design. This makes serial links such as FireWire more attractive for commodity links.

4.3. Magnetic recording trends

Since 1991, magnetic disk drive areal densities have been advancing with a rate of progress of 60% per year. It is believed that this rate of progress can be sustained at least for another 5 years and perhaps even increased within this time frame. However, there are some limits that will be reached in the near future that may require significant change in this technology. Since its introduction, the areal density of magnetic disk recording has been increased over 2 million times by linear scaling of dimensions of the head, medium and head-medium spacing, while the sensitivity of read heads has been increased. Since the signal to noise ratio scales approximately with the number of magnetic particles contained within a bit, it has been necessary to reduce the size of the magnetic particles in order to advance the performance forward. Ultimately, the particle size will become so small that the magnetic energy of a particle will decrease to values that approach thermal energy. When this occurs, thermal energy alone may cause magnetic recordings to become unstable.

The magnetic energy of a particle is given by $K_U V$, (where K_U is the magnetic anisotropy energy holding the magnetization in its orientation and V is the volume of the particle) while the thermal energy is $K_B T$ (where K_B is Boltzmann's constant and T is the absolute temperature). Thus for reliable magnetic storage the ratio $K_U V / K_B T$ must exceed a certain threshold. Magnetic materials, which have higher values of magnetic anisotropy energy K_U , exist; however, they require high fields for recording. Presently such fields cannot be produced with existing recording heads. The heads saturate before such fields can be produced.


Modeling has indicated that if magnetic recording continues to be scaled linearly, this thermal instability limit will be reached at about 36 Gbit/in². At the 60%/year growth rate, this density would be reached in about 5 years. However, modeling performed by the National Storage Industry Consortium has recently shown that, if the bit cell is not scaled linearly, but decreased in track width more than in bit length, a density of 100 Gbit/in² may be achievable. Ultimately, however, thermal instabilities will become a roadblock. One possible means of extending magnetic recording density beyond what can be achieved using longitudinal magnetic recording is to use perpendicular magnetic recording where the medium is magnetized perpendicular to the film plane, rather than in the plane. If a high permeability magnetic underlayer is placed under the perpendicularly magnetized thin film medium, then an image of the magnetic head pole is produced in the underlayer. Consequently, the thin film recording medium is effectively kept within the gap of a recording head producing larger recording fields. With this larger record field it is possible to record on media with higher K_U and, consequently, smaller grain size V and smaller bit sizes. Although this approach has been investigated for more than two decades it has not been able to compete with longitudinal recording because of higher cost. This suggests that in the future it might be difficult for magnetic disk drive industry to maintain the cost/performance figures within present trends. Another possible means to circumvent the thermal stability problem of conventional longitudinal magnetic recording is to use thermally

assisted recording process. If the media is kept near room temperature for storage and readout, but raised in temperature during recording it may be possible to use media with intrinsically higher K_U and, consequently, improved thermal stability while still being able to record on them. One way of introducing heat is by means of a focused laser beam, which leads to magneto-optical recording. This direction has resulted in a renewed interest on this optical storage approach.

5. Advanced Optical Storage Technologies

5.1. Parallel Access

With the strong demand for capacity for removable systems comes an equally strong demand for high data rates. This demand originates from the wish to quickly transfer large image and video files to direct memory and from the desire to perform fast content based search and image processing with this type of files. The conventional method to increase data rates in a disk system is to increase the rotation speed and to increase the linear bit density. Present rotation speeds are already high and are limited by media integrity and servo speeds. The linear density increases as the square root of the areal density and therefore falls short of satisfying the emerging needs of content based database search.



SOURCE: Fujitsu

Figure 26. Multi-beam optical head investigated at Fujitsu.

(Source: Fujitsu)

Based on progress made in different areas of optoelectronics, newer technologies exploiting the parallel access capabilities of optical storage are emerging, to satisfy these new requirements. In Japan, short term and long term approaches to parallel accessing data on optical disks are under investigation. For example, Fujitsu is investigating the use of laser and detector arrays to access data in parallel from an MO disk as shown in Figure 26.

Issues including heat extraction from laser diode arrays, multi-beam positioning errors, packaging and signal processing circuit complexity for post processing are being investigated.

Zen Research³⁹ was founded in 1994, and has subsidiaries located in Cupertino, CA, and Tel Aviv, Israel. Zen Research sells Zen TrueX/Multibeam components to optical drive manufacturers to incorporate into their products. The Zen TrueX/Multibeam component sets include optics, detection devices and high speed ASICs which enable drives to illuminate

multiple tracks simultaneously, read them in parallel, and process the data streams through a custom ASIC, in a manner compatible with all existing CD and DVD formats and standards. Zen has partnered with a number of drive manufacturers who sell to PC makers. Currently, Kenwood Technologies Inc., a subsidiary of Kenwood Corporation, offers two Zen-enabled CD-ROMs (the MULTI BEAM 40X Plus™ and the 52X TrueX™, see Figure 27), and plans to increase their offering to include a Zen-enabled DVD in 1999. Zen Research's vision includes 60X CD-ROM and 12X DVD by 1999. The company has been awarded 10 patents on this core technology, with another 17 pending or in process.




Figure 27. Kenwood's \$99 MULTI BEAM 40X Plus™ CD-ROM drive (source: http://www.kenwoodtech.com/home_toplink.html)

CD-ROM prices have consistently decreased while performance has only marginally increased, leveling off at 32X. Users typically see much lower rates, since conventional CAV (Constant Angular Velocity) drives, which comprise the bulk of models on the market today, allow for a maximum speed of 32X or higher, but only on the outermost tracks. Data read from the inner tracks of a disk, where most of today's software is located, is read at only 12-16X (see Figure




Figure 28. Sustained data transfer rate comparisons for serial CD and DVD -ROM and 6-channel parallel CD-ROM drives (source: http://www.zenresearch.com/product_info4.html).

28). Zen claims that TrueX enabled CD-ROM drives will extend the usefulness of CD-ROM

technology for many years. Zen will also apply TrueX technology to increase the performance of DVD products.




Figure 29. Multibeam optical head schematic for Zen CD-ROM drives (source: http://www.zenresearch.com/product_info4.html).

Zen's approach splits a laser beam into seven beams to read seven tracks in parallel, as in Figure 29. The center beam is used for focus and tracking, while three beams on either side read the multiple tracks. The data is then processed through a custom, single-chip ASIC for parallel processing and error correction via an integrated signal processor. A split head is used to reduce weight on the movable portion, in turn lowering access times. By varying the speed the disc spins (Zen currently uses a CLV approach), Zen achieves over 6MB/s of data rate throughout the entire disc radius. Spindle vibration problems are nonexistent because the discs spin at only 3,000RPM, roughly the speed of a 6x to 10x drive.




Figure 30. Schematic of a VCSEL array access with a massively parallel near field optics architecture to optical disks promising TB capacities and Tb/s data transfer.

A more futuristic approach being explored at the Tokai University by Prof. K. Goto and his group is shown in Figure 30⁴⁰. This approach involves the use of VCSEL arrays flying directly above an optical disk medium using a near field geometry. The laser is used to both record and directly detect the presence of a bit based on the optical signal that is fed back into the laser. In order to achieve high data rates, massively parallel access through VCSEL array is contemplated, using VCSELs as described in Figure 31.




Figure 31. Use of a VCSEL for recording and readout in a near field optics geometry


(Source: K. Goto, Tokai University)

In the U.S., Quinta Inc.⁴¹ was recently acquired by Seagate Inc., a manufacturer of magnetic hard drives, has been developing a method to access several optical disk surfaces in parallel using optical fibers. The fibers are used for distributing the laser power to the head and back to the receivers for readout and a novel optoelectronic switch is used to selectively access data on various plates. At COMDEX/Fall '98, Quinta demonstrated a 5.25", half-height, 6 platter, 12 head disc drive, running at 4500 RPM with a 10 msec seek time, with an areal density capability of 3-4 Gbits/in², the equivalent of 4.0 to 5.4 Gbytes per platter for a 3.5" disc drive. In November of 1998, they announced a joint agreement with Mitsubishi Chemical Corporation (MCC) to co-develop advanced recording media for Quinta's optically assisted Winchester (OAWTM) architecture.

In the U.S., Call/Recall Inc. is leading the FROST consortium (see section 0) in investigating the use of VCSEL Arrays and MEMS to facilitate parallel access to multilayer disks. In addition, massively parallel accessing of holographic data storage systems is also under way as described in section 5.3. It is believed that parallel access is a viable method to increase the data transfer rate of optical and probe storage systems. The issue is rather when the necessary components that can enable massively parallel data transfer may provide a favorable cost entry point to market such systems.

5.2. Near-field optical recording (Terastor)

Terastor's near field optical recording (NFR) approach, shown in Figure 32, is based on work by Gordon Kino at Stanford University done in the early 1990s.^{42,43,44} When a solid lens shaped like a truncated sphere, a solid immersion lens (SIL), is placed between a standard objective lens and the media surface, the SIL focuses the incident rays of the laser from the objective lens to a single spot at the base of the partial sphere. The light at the outside edges of the focussed beam is focussed at very high angles of incidence to the flat focal surface. This outer ring of the beam is totally internally reflected within the SIL unless the SIL is brought to within a fraction of a wavelength of the media. At this close proximity, a portion of the normally reflected light will couple evanescently through the SIL to the media, resulting in a spot approximately half the size of the spot achieved by using only an objective lens. Maintaining such a distance over a rapidly spinning disc is possible using a technique borrowed from hard disk storage: the flying head. The flying head floats on a cushion of air above the disc and is able to hold the head in very close proximity to the disc without crashing. SIL's have been fabricated with equivalent numerical aperture of up to 1.8. Using 780nm wavelength, recording spot sizes on the order of 320 nm have been achieved giving recording densities on the order of 3 Gbit/in². By modulating the magnetic field appropriately, only a "crescent" of this spot is recorded, to further increase the linear density. Using the flying head, data rates of up to 3 Mbit/sec were achieved and data rates of up to 15 Mbit/sec may be possible.



- A** Laser Beam
B Objective Lens
C Flying Head
D Solid Immersion Lens (SIL)
E Plastic Substrate
F Evanescent Coupling
G First Surface Recording
H 6 μm fly height
I Magnetic Coil in the head

source: Terastor




Figure 32. Terastor's near-field optical approach using MO media, and their projected areal density roadmap comparison.

(Source: Terastor)

Terastor's roadmap is shown in Figure 32. The first 20 GB removable NFR disk products were initially due to be offered in late 1997, but this delivery was postponed to late 1998, and now the first product is planned to be a 10-Gbyte drive in early 1999. The \$700 to \$800 drives will have a sustained data transfer rate of 6 Mbytes per second. A 20-Gbyte drive, due in the second quarter of 1999 for \$1,000 to \$1,200, will offer an 11-Mbyte-per-second sustained transfer rate. Both will have an 18-millisecond seek time. The first drives will use 5.25-inch, single-sided removable media, though dual-sided media is planned to be available later. Quantum Corp. has acquired a license to make the removable devices. The reasons for the delays have not been announced, however, recent analysis by Prof. Mansuripur at the Optical Data Storage Center at the University of Arizona has indicated that the evanescent coupling needed to achieve the reduced spot size will necessitate fly height of less than 1 micro inch (25 nm).⁴⁵ This will likely result in either lower areal densities or non-removable media, or both.

Terastor's 10 GB drives are projected to immediately compete with removable HDD and magneto-optical in applications such as medical imaging, desktop publishing, CAD/CAM,

multimedia and graphic design. Their 20 GB drives are targeted towards those markets as well as traditional enterprise online backup and archive applications.

5.3. Volume holography

Holographic storage for data storage has been investigated since the late 1960's, largely due to its theoretical potential for offering high capacity, high transfer rate, and low access times. Despite the periodically intense interest in this technology over the last 30 years, progress in realizing commercially viable systems has been limited by media and system issues. These issues are again being seriously investigated, at universities, small and large companies.⁴⁶ This technology received significant DARPA funding in the 1990's via the Photorefractive Information Storage Materials (PRISM) and Holographic Data Storage System (HDSS)




Figure 33. Principles of 90 degree geometry holographic recording and reading

(Source: Optitek, Inc.)

programs. The five-year, \$32 million HDSS program is being supported 50 percent by the U.S. Defense Department's Advanced Research Project Agency (DARPA) and 50 percent by the 12 participants. The HDSS and PRISM programs both involve many of the same participants, have the same principal technical investigators, and are funded by DARPA and the participants. The PRISM project is developing optically sensitive materials optimized for storing holograms and an understanding of the various tradeoffs that must be made between mutually exclusive performance parameters. HDSS is developing the other hardware technologies needed for practical holographic data storage systems and integrating them into demonstration systems. The HDSS program was formed demonstrate write-once and rewritable systems with a capacity of 1 trillion bits or more and a data-throughput rate of at least 1 billion bits a second. The HDSS participants are Stanford University, Carnegie-Mellon University, GTE Corp., IBM's Almaden and Watson Research Center, Kodak, Optitek, Rochester Photonics, Rockwell, SDL, Inc., University of Arizona, and University of Dayton (Ohio). The National Storage Industry

Consortium (NSIC) manages the financial and administrative aspects of the HDSS program. Almost all of the recent holographic storage demonstration experiments have been produced by HDSS members, or by holographic storage groups at Caltech and Holoplex, Inc..

To create holograms, the interference fringe pattern of two optical wavefronts creates holograms, as in Figure 33. The storage media can record the fringes as index and/or amplitude modulation. When the recording is illuminated by one of the wavefronts (the reference beam), the other wavefront (the object beam) is reproduced. The recording of 2D data pages can enable parallel reading for high data transfer rates. Bragg selectivity facilitates the storage of multiple pages within the same volume and retrieval via presentation of the appropriately coded reference beam.⁴⁷ The most significant feature of holographic storage is its potential to offer very fast access to an arbitrary data page. This capability is facilitated through the use of massless addressing or "multiplexing" using wavelength, angle, phase code, etc. In order to further increase storage capacity, massless addressing is generally combined with other multiplexing methods such as peristrophic⁴⁸ (media rotation), shift⁴⁹ (fine translation of the storage medium), and space multiplexing (gross translation of the holographic recording location). Combining these last three cases can achieve very high capacities; however, some mass is required to move before access to certain data pages can be obtained.

However, for practical (i.e., moderate power) laser sources, the time required to accumulate sufficient signal photons over the retrieved data page determines the access time of the memory system. Since the number of photons associated with a retrieved bit per unit time, per unit readout power is inversely proportional to the square of the number of multiplexed pages, there is once again a direct trade-off between capacity and access time.⁵⁰ This trade-off is governed not by the mechanical system associated with motion control, but rather by the constant of proportionality between the holographic diffraction efficiency and the inverse of the square of the number of stored pages. A significant materials research effort is underway to develop recording media for which this constant is large. There have been some attempts to describe holographic storage systems in the context of a high capacity DRAM replacement⁵¹; however, such a solution is extremely cost sensitive and competes with a rapidly moving target. It is fair to say that most holographic storage systems will be marketed as high capacity (> 10 GB) secondary storage with near-term focus on low-cost media distribution ROM. Both WORM and R/W systems are feasible within the holographic paradigm; however, materials limitations and the relative immaturity of *in situ* hologram fixing techniques make ROM an attractive first market option.⁵²

Currently, photorefractive crystals (PRC) are typically used in non-moving media holographic demonstration experiments, and photo-polymers are used in volumetric disk media demonstrations. Both of these materials are recorded at room temperature and the inorganic PRCs offer high photocyclicity. The main drawback of the approach with PRCs, is the high cost of materials, the decrease in diffraction efficiency with increasing number of multiplexed holograms and read-out cycles, and the relatively small and slow changes in the index of refraction. To circumvent these problems, significant progress have been made in terms of using

recording schedules to improve the uniformity of the recorded data, fixing procedures to permanently store data in PRCs and new recording geometries to minimize crosstalk. In addition, more recently, photorefractive polymers that exhibit large index of refraction changes under very large electric fields have been demonstrated. Holographic storage as a read only archival memory can also use polymers such as available from DuPont and Polaroid. In this case, the shrinkage of the polymer after recording and over time results in a shift in readout wavelength, as well as in increased crosstalk and loss of resolution. Very recently encouraging results on reduced shrinkage materials have been announced, but the dimensional stability of these new materials still appears to be an order of magnitude smaller than what is needed for commercialization. Another issue in removable media holographic systems will be the tight wavelength tolerances that must be met between the recording and replay systems. This will likely require reading lasers that can be tuned to exactly match the recording wavelength, and may prevent the cost reductions that have made CD-ROM systems so ubiquitous.

Recent experimental system demonstrations have made impressive progress, and a summary of both non-moving media and moving media demonstrations is given in the next paragraphs. This summary is followed by a discussion of the future potential of such systems, and the theoretical and the practical roadblocks limiting their future performance.

The storage capacity associated with single location recording is determined by many factors that depend on both material and system variables. Perhaps the most important factor is the so-called $M/\#$ of the system.⁵³ The system $M/\#$ is defined with respect to the readout diffraction efficiency η as a function of the number of multiplexed pages M : $\eta(M) = (M/\# / M)^2$. $M/\#$ can also be written in terms of the holographic recording and erase time constants τ_r and τ_e , as $M/\# = A_0 \tau_e / \tau_r$ where A_0 measures the saturation diffraction efficiency of the system. This parameter is determined primarily by the recording material and reasonably large $M/\#$ has been measured in recent photopolymer systems.⁵⁴ The primary drawback associated with the use of photopolymer is its limited thickness (generally $<100\mu\text{m}$). These thin samples can not achieve high capacities without physical motion. For this reason, the largest capacities for single position storage have been achieved using photorefractive crystals. The example system described here is based on LiNbO_3 . We should note that recent successes in the creation of thick polymer samples together with novel laminating techniques suggest that single spot capacities comparable to those described below will be achieved in these systems very soon.⁵⁵

Geoff Burr at Caltech has reported the storage of 10,000 holograms at a single position in photorefractive lithium niobate.⁵⁶ His system utilized 4 fractal rows (vertical angular addressing) and 2500 angles (horizontal angular addressing). The required angular separation was 1.4 degrees vertical (image bandwidth was 1.15 degrees) and 0.004 degrees horizontal. This experiment used a mechanical scanner but an acousto-optic (AO) scanner could be used to offer massless non-mechanical addressing. Each data page contained 480×440 pixels and the extension to 1000×1000 pixels offers no conceptual difficulty in this architecture. Although SNR measurements were reported, no coding was utilized and little further attention to data fidelity

was paid. Fresnel plane storage was used to mitigate photovoltaic effects; however, this will not be an issue in photopolymer systems or alternate photorefractive such as SBN. These 10,000 holograms were stored in a $1 \times 1 \times 2 \text{ cm}^3$ crystal and the measured capacity was 2.11 Gbits. This corresponds to a volumetric storage density of roughly 1 Gbit per cm^3 . More relevant perhaps is the areal storage density which, for the object beam dimensions in this demonstration (1.77mm by 2.4mm), was nearly 500 bits per μm^2 . The average diffraction efficiency was measured to be 5×10^{-9} setting an ultimate limit on data fidelity and transfer rate. AO addressing could achieve access times in the range of 10-100 μs and acceptable SNR was achieved at video rates, corresponding to a data transfer rate of 6.4 Mbps.

Other single location demonstrations are notable. These include the 3D disk-based system described by Li⁵⁷ in which he examines geometric constraints. He does not consider material limitations and analyzes both angle and wavelength multiplexing to find optimized surface storage densities. He predicts that wavelength MUX using a 3cm thick sample can achieve 166 bits per μm^2 . Pu⁵⁸ examines peristrophic multiplexing in thin materials (100 μm photopolymer). He demonstrates 10 bits per μm^2 and predicts that > 100 bits per μm^2 is possible. Others have considered various multiplexing methods to increase the practical capacity and/or density; however, no demonstrations have measured capacities in excess of the Burr experiment. For example Campbell⁵⁹ and Hesselink⁶⁰ have discussed combining angle and wavelength multiplexing. Although such a scheme does not offer an inherent capacity advantage (because the same volume of k-space is available using wavelength multiplexing alone), there is a practical advantage associated with combining >1 multiplexing dimension due to the reduced requirements associated with each independent scanning apparatus (e.g., combining a small degree of source frequency agility with a relatively modest number of angles). Hesselink, Walkup, and others have also considered the use of phase code multiplexing for similar reasons; however, sensitivity issues concerning phase-repeatability severely limit this approach.^{61,62,63}

The Rockwell holographic storage demonstrator is a relatively compact system with dimensions of roughly 9"x 12"x 5". It is based on AO devices and has been demonstrated at full capacity with nonmechanical angle multiplexing offering access time in the range of 10-100 μs . Data rate in this system remains limited by CCD frame rates. The Rockwell geometry utilizes multiple object beam locations along the path of the reference beam. While recording in this geometry requires physical motion of the media within a docking station, readout requires adjustment of the reference angle along one-axis only. An output shutter selects the desired reconstruction. Little has been published regarding the detailed specifications of this system. The capacity of the system is a few Gbytes; however, the demonstration is important because it has successfully undergone shock, environmental chamber (temperature and humidity) cycling, and lifetime testing for airborne applications.

Multiple locations may also be used to increase the system capacity, with the consequence of slower access time. Linear motion of rectangular solids (e.g., photorefractive crystals) has been proposed in combination with angle and/or wavelength multiplexing to achieve high capacity.⁶⁴

From a mechanical perspective however, disk-based systems are more attractive than those using linear motion and there is the added benefit of utilizing some existing infrastructure in such holographic disk-based systems. Disk-based systems using photorefractives have been




Figure 34. Holographic demonstration systems constructed at Rockwell Sciences Center.⁶⁹

described, that combine angle multiplexing or peristrophic multiplexing in a single location with radial (space MUX) and rotational motion (shift MUX).⁶⁵ These systems are costly and complex due to the need for mechanical servos over many additional degrees of freedom (radial, reference angles (2), disk tilt, etc.). Disk-based architectures utilizing thin materials (e.g., photopolymers) have been impressive in achieving high surface storage densities via shift multiplexing.⁶⁶ The use of polymers has also reduced media cost in these systems. The use of inexpensive materials together with recent demonstrations of low-cost/high speed replication has provided a near-term potential for mass distribution holographic ROM.⁶⁷ Li⁶⁸ has examined disk-based systems using angle multiplexing and wavelength multiplexing. Each location contained some number of multiplexed holograms and conventional means were envisioned to access different storage locations on the disk. His results suggest that storage density is limited by geometric factors related to the scanners and NA of the optical systems. Optimum designs for angle multiplexing can yield 120 bits per μm^2 , and optimum wavelength multiplexed designs are predicted to offer 166 bits per μm^2 . These densities correspond to single unit capacities of roughly 1 Tbit in a form factor similar to conventional optical disk drives.

Thin holographic materials can use shift multiplexing to achieve high surface densities. Shift multiplexing is based on non-planar reference beams. Demonstration systems have used spherical waves, multiple plane waves, and more recently speckle patterns.⁷⁰ Shift multiplexing is attractive in disk-based systems since only radial servos are required for addressing. This makes the mechanical aspects of such a system similar to its bit-wise counterpart; although, disk tilt remains a critical issue in these systems and may represent a significant cost/reliability factor. Predictions for areal densities of 20 bits per μm^2 have been made. Demonstrations using 100x100 bit pages in 38 μm thick Dupont samples have measured shift selectivities of 2.8 μm and corresponding densities of 10 bits per μm^2 . The predicted density of 20 bits per μm^2

corresponds to a single unit capacity of 25GB without significant additional cost as compared with existing optical disk drives.

In the next paragraphs, the capabilities of future holographic storage systems are predicted, and potential roadblocks are discussed, beginning with estimates of the performance of single location systems. There are three classes of roadblocks associated with the performance of these systems. The first of these relates to the capabilities of holographic materials and the impact that material characteristics will have on the capacity/access time trade-off. The maximum number of stored pages (M) is related to the access time (T), readout laser power (P), and system $M/\#$ according to $M = M/\# \sqrt{T} \sqrt{P/K} N$, where N is the number of pixels per page and K is a constant determined by data fidelity requirements. This expression offers an optimistic estimate since in reality, K will be a function of T . Requiring 100 signal photons for reliable detection yields $K=4 \times 10^{-17}$ J. Taking typical values for $N=10^6$ pixels, $P=1$ W, and $M/\#=1$, we obtain a useful rule of thumb: $M=5000 \sqrt{T}$ for T in ms. This provides a 5 Gbit capacity with 1ms access time. This performance can be achieved using any nonmechanical multiplexing technique such as angle, wavelength, peristrophic, etc. It is useful to note that the Burr angle MUX demonstration system achieves nearly this material limited performance. This simple analysis also provides a means for estimating future performance. Taking higher source power $P=10$ W and larger $M/\#=10$ we obtain $M=1.6 \times 10^5 \sqrt{T}$, yielding a more useful 20 GByte capacity and 1ms access time. It may be argued that an important advantage of the single location architecture is its potential for very fast random access; however, the capacity cost associated with utilizing this capability is clear. For the 20 Gbyte system described above, reducing the detector integration time to AO addressing times of $T=10 \mu$ s, reduces the system capacity to 2 Gbyte.

The above analysis concerns material limitations only inasmuch as they directly impact diffraction efficiency. It was implicitly assumed that sufficient area and thickness could be achieved. It is also important however, to consider additional material properties such as homogeneity and surface quality (to control scattering), mechanical and chemical stability (to insure acceptable environmental performance), and cost. Polymer materials have a potential advantage along both the diffraction efficiency and cost dimensions; however, there remain problems concerning thickness, optical quality, and stability in these material systems.

The second class of roadblocks associated with the future performance of holographic storage systems relates to the capabilities of supporting technologies such as lasers, detectors, scanners, etc. The above example assumes the existence of a 10W laser whose characteristics (i.e., wavelength, linewidth, stability, compactness, cost, etc.) are suited to a commercial storage application. It also assumes the existence of a multiplexing scheme (or combination of schemes) capable of addressing 158,000 pages in a single location. Existing AO based solutions are capable of implementing such an architecture (e.g., 2,000 horizontal angles and 79 vertical angles); however, the cost of this implementation may be prohibitive, requiring two AO cells and compensating AO or EO devices. System volume will also become a concern with an AO based solution. The existence of compact, low cost, frequency agile lasers would therefore

significantly improve the commercial success of the single location system by facilitating wavelength multiplexed solutions.

Another critical device technology for the realization of holographic memories is the 2D detector array used for interfacing to the optical system. This device determines the constant K , in the capacity/access time trade-off and, via its readout channel, determines the sustained data transfer rate as well. CCD cameras with 1024×1024 pixels and frame rates of 1 kHz are now available as research class devices.⁷¹ Such a device will be commercially mature in 2010 and can offer transfer rates of 1 Gbps; however, the serial nature of the readout channel is suboptimal for the storage application. 2D CMOS detector arrays will offer a lower power, more compact, low cost user interface solution. The potential for (1) random access to an arbitrary pixel, (2) integrated parallel channel electronics and (3) active tracking/timing circuits, make the CMOS array solution attractive; however, the sensitivity and noise characteristics of these devices must be improved prior to their insertion into high capacity holographic systems. For example, we can refine the capacity estimate from above by including the detector noise characteristics more carefully. If the detector noise is given by a noise equivalent power of 4.5×10^{-15} W/ $\sqrt{\text{Hz}}$ and a $P=1\text{W}$ readout power is used, then a 1 ms access time corresponds to an information theory limited capacity of 2 Gbits.⁷² Once again increased laser power $P=10\text{W}$ and increased $M/\# = 10$ can produce larger capacities by roughly a factor of $10 \times \sqrt{10} = 32$ producing $C=8$ Gbytes.

The third class of roadblocks associated with the future performance of holographic storage systems relates to considerations of reliability, cost, environmental sensitivity, removability, ease of replication, media volume, system volume, etc. An understanding of these issues is required to establish the marketability of the holographic storage alternative. Unfortunately, little can be said concerning these "engineering" considerations owing to the lack of a performance history along any of these dimensions. The Rockwell demonstration has begun to test the waters of military specifications and environmental tolerancing, but no firm conclusions can be drawn as yet. Some studies of system volume have been performed for the single position architecture based on a 4f optical system.⁷³ Given an input pixel size and system $M/\#$, it is possible to design an optical 4f system that balances the field size (number of pixels per page) against the required media volume to maximize the media storage density. An input pixel size of $15 \mu\text{m}$ for example, produces an optimal focal length of 21.5mm for 1 Megapixel pages in the paraxial limit. Using $M/\#=1$, the results of the optimization in this case suggest that a 1 Gbit capacity can be achieved in a material thickness of 4.7mm with a cross-section of 8.25mm^2 , producing volumetric media densities of roughly 0.025 bits per μm^3 and areal densities of 121 bits per μm^2 . The system volume in this case is dominated by the required lens spacing and is found to be roughly 87.6cm^3 . It is also possible to minimize the system volume via selection of the most compact optical design. In the case cited above this produces a reduction in system volume to 6.1cm^3 and a corresponding reduction in media density by a factor of roughly 100 both volumetric and areal. Although these results are preliminary and a similar analysis must be applied to systems using random phase masks and other non-4F

architectures, it is interesting to compare this result with the often quoted V / λ^3 result of roughly 1 bit per μm^3 .

The capabilities of holographic storage systems based on spatial multiplexing without physical motion will arise primarily from the limitations of scanning devices. The Burr system can be taken as a design baseline. We might imagine single position multiplexing using AO horizontal angle (2000 angles) and wavelength converted to vertical angle for fractal multiplexing (79 wavelengths). Suppose that all of the single location capacity possible can be exploited using these two degrees of freedom. A segmented mirror array can be used to convert the angular output of an orthogonal AO device to vertical position without requiring any physical motion. Once again taking the number of resolvable angles to be 2000, we achieve a capacity of $2000 \times 8 \text{ Gbyte} = 16 \text{ Tbyte}$ with a photon-count limited access time of 1ms. Exploiting the nonmechanical nature of this system by reducing the detector integration time and to take better advantage of the AO addressing once again incurs a capacity cost. Taking $T=10 \mu\text{s}$ yields a capacity of 1.6 Tbyte.

However, the material and system volumes required to achieve this capacity make such a system impractical. The spatial multiplexing will utilize 2000 times the material height of the single position approach. The Burr system uses a vertical separation of 3mm resulting in a required 6m of material! Such a system is clearly unrealistic. We can estimate the minimum vertical spacing Δ_v as a function of various optical system parameters by requiring that the spatial multiplexing distance exceed the extent of the data Fourier transform along this dimension. That is $\Delta_v = \lambda f / d$, where d is the data pixel size and f is the lens focal length. For realistic values of $f=10\text{cm}$, and $d=20 \mu\text{m}$, we find $\Delta_v = 0.25 \text{ cm}$. Even considering a rather large crystal height of 5cm, it is only possible to multiplex $50/0.25 = 20$ vertical positions. Such a system offers a capacity of 160 Gbyte with 1ms access time.

The projections of the previous paragraphs depend upon improvements in material ($M/\# = 10$), laser ($P=10\text{W}$ with 79 wavelength tunability), detector (acceptable fidelity at 100 photons), and scanner technologies. Relaxing these projections will result in a system comparable to the Burr demonstration of $< 100 \text{ Gbits}$ capacity. The geometry of such an architecture is not well suited to a disk-based implementation; however, linear motion can be added in a simple way to further increase capacity with the concomitant cost in access time associated with moving the large material mass required for this geometry. The need for thick materials, four AO/EO scanners, and multiple 4f optical systems makes this approach potentially unviable from a commercial perspective. It is therefore desirable to examine the potential of holographic systems based on a more cost effective architecture.

Many recent studies of disk based holographic systems have focused on the optimization of effective surface storage density. This is due in part to the desire to use available materials in convenient thickness. The study by Li⁷⁴ has predicted an optimized density of 166 bits per μm^2 . This result is based not on material limitations or fidelity requirements, but on recording/readout geometry together with tunable source capability. This result is therefore a

good news/bad news story in which 1.4 Tbit = 175 Gbyte can be stored in a single media unit with dimensions comparable to conventional optical disks. It is not unreasonable to expect system form factors for wavelength multiplexed holographic disks similar to those found in today's optical disk drives. The bad news however, is that improved material characteristics do not directly impact this capacity. The primary technological dimensions along which this performance can be improved concerns source technology. In particular, the storage density is proportional to the range of wavelength tuning $\Delta\lambda$ and inversely proportional to the cube of the minimum achievable wavelength λ_1 . The above estimate of 166 bits per μm^2 was made for $\lambda_1=500\text{nm}$ and $\Delta\lambda = 40\text{nm}$. Reducing λ_1 to 450nm and increasing $\Delta\lambda$ to 100nm would increase the density to 570 bits per μm^2 . This density yields a single media unit capacity of 600 Gbyte.

A fine point regarding these disk based systems concerns the precise method of readout during disk rotation. Acceptable signal strength requires 1ms of detector integration time; however, continuous disk motion would prohibit such an extended illumination of a single recording location. An alternate readout scheme using a pulsed source might alleviate this problem, except that very high peak powers will result. For example, if the mechanical dwell time is only 1 μs , then the previous 10W power requirement becomes a 10kW peak power requirement in the pulsed case. An alternate stop/start rotation scheme might be utilized; however, the access time of such a system will not compete with conventional continuously rotating solutions.

The previous discussion concerns predicted performance of holographic ROM systems. The photorefractive materials on which most of that discussion was based, do not offer sufficient recording sensitivity to realize R/W or WORM capability. Projections for R/W and WORM systems are critically dependent upon the development of sensitive recording materials and recent photopolymers have shown significant capability in this regard. Sensitivities greater than $0.1 \text{ cm}^2 / \text{mJ}$ and refractive index modulations greater than 1×10^{-2} have been measured in materials that can be used in thickness up to 250 μm .^{75,76,77} These sensitivities suggest that a 1W source can be used to produce reasonable holographic recordings with 1ms exposure times. Additional effort is required to understand issues related to pre-exposure, shrinkage, local "fixing," etc.

Table 3. Demonstrated and potential capacities and densities for volume holographic data storage.

Demonstration			Forecast		
Capacity	Volumetric Density	Areal Density	Capacity	Volumetric Density	Areal Density
Motionless Media: 1 Location			comment: .25x.25x4 cm ³		
2Gbit	1Gbit/cm ³	500bit/μm ²	8Gbyte	32Gbyte/cm ³	1.3kbit/μm ²
Motionless Media: > 1 Location			comment: 5x.25x4 cm ³		
34Gbit	40Mbit/cm ³	236bit/μm ²	160Gbyte	32Gbyte/cm ³	1.3kbit/μm ²
Rotating Disk			Comment Δλ=100nm		
100Gbit	(t=38 μm)	10bit/μm ²	600Gbyte	=NA	570bit/μm ²

* Predictions based on P=10W, M/#=10, and 1ms access time

Projected applications for volume holographic storage include large-scale archiving and library databases, and the long-term potential for optically implemented content-based addressing might add value in these applications. Recently, the use of photorefractive volume holography to implement a DRAM replacement has also been suggested.⁷⁸

5.4. *Persistent spectral hole burning / temporal holography*⁷⁹

In another non-moving media approach, Persistent Spectral Hole Burning materials enable wavelength to be used as the third dimension. This spectral memory is investigated at the University of Oregon, and has spawned a small company, Templex Technologies, Inc..⁸⁰ As shown in Figure 35, single bits within a data stream are stored throughout a wide bandwidth and spectrally overlap bits earlier and later in the bit sequence.⁸¹ In ordinary spatial holography, interference between two light fields can lead to the storage of a laser beam's spatial wavefront information. The stored information allows for a beam's complete reproduction. In spectral holography, two finite duration beams (simultaneous or not) interact with a frequency-selective recording material. Interference of the two beams in frequency space leads to the storage of one beam's temporal waveform information (represented by a frequency-dependent absorption). If the optical beam is encoded with data, that information is included in the recorded waveform. Readout of spectral holograms produces a signal beam whose temporal profile duplicates the original input data beam.⁸² Since frequency-selective storage materials are also spatially selective, it is possible to make spatial-spectral holograms in which both the temporal and spatial structure of input beams are recorded.

Spectral memory demonstration experiments at the University of Oregon employed a 780-nm commercial semiconductor diode laser as the light source, a crystal of Tm^{3+} :YAG as the frequency-selective recording material, and an avalanche photodiode as a signal detector. The diode laser was stabilized to an external cavity containing a grating and an electrooptic crystal. The intracavity electrooptic crystal provides for microsecond-time-scale sweeping of the laser frequency over roughly one gigahertz. Two storage (reference and data) beams and one reading beam, are created from the output of the single laser source using the beam splitter and the acousto-optic modulators. The beams are focussed to a 150 m^2 spot in a Tm^{3+} :YAG crystal. The reference and data beams are simultaneous as are the read and signal beams.

A 1760-bit-long sequence of 20 Megabit/sec data was encoded on an input beam, stored, and recalled. The areal density achieved is about 8 Gbit/in^2 and the density-bandwidth product is $1.5 \times 10^{17} \text{ bits/in}^2\text{-sec}$. In the demonstration material, areal densities of up to 100 Gbit/in^2 and bit rates up to 1 Gbit/sec appear feasible. To achieve these results in Tm^{3+} :YAG (useable with convenient 780-nm diode lasers), it was necessary to work at cryogenic temperatures ($< 10 \text{ K}$) and tolerate tens of millisecond-scale storage persistence. To date, however, an ideal storage material has not been identified. Memory is but one application of spectral-spatial holography. Temporal waveform processing, temporal-waveform-controlled optical beam routing, and optical bit-rate conversion have all been demonstrated. These processes are useful in communications and other applications and in many respects place looser demands on materials than the memory application does.

Templex Technology, Inc. was commercializing this approach to optical data storage, under the name of Temporally Accessed Spectral Multiplexing (TASM). The company has apparently now decided to pursue the communication applications mentioned above via an implementation of an optical Code Division Multiple Access (CDMA) communications systems. They anticipate their CDMA optical communication system will provide for the multiplexing of multiple high bandwidth communication channels to achieve Terabit per second data rates over a single fiber.

The currently limited capacity and lifetime, but potentially high transfer rates of this storage medium may lead to applications as high-speed cache for processors, though the first proposed commercial application was at an intermediate level in the data storage hierarchy between solid state memories (DRAM/SRAM) and magnetic hard disk storage. The potential for processing (matching, filtering) the data as it is retrieved may also lead to "smart memory" applications, as well as the telecommunications applications mentioned above.




Figure 35. Temporal holography using persistent spectral hole burning
(Source: Templex)

5.5. Pit Depth Modulation

Another degree of freedom that is being considered Calimetrix, Inc.⁸³ is the recording multiple values per spot, by varying the depth of pits within the media. The depth of the data pit varies, while the length is fixed, as in Figure 36. The photocell detects the variable pit depths and produces a graded, rather than binary, output. Pit-depth modulated disks are projected to provide 15 GB of data at rates of 19.4 Mbit/s for HDTV films, using current red laser (635nm/650 nm) diodes.

Calimetrics has also been awarded an Advanced Technology Program (ATP) grant from the National Institute of Standards and Technology to further develop its technology for writable optical storage systems. Calimetrics is the lead company in a joint venture administered by the National Storage Industry Consortium and including Polaroid Corporation, Energy Conversion Devices, Inc., the University of Arizona's Optical Sciences Center, and Georgia Institute of Technology. The project, known as the Multiple Optical Recording Enhancements (MORE)




Figure 36. Pit depth modulation: principle and result
(Source: Calimetrics)

Joint Venture, expects to increase today's rewritable and write-once DVD capacity (currently 2.6GB to 3.9GB per disc, expected to move to 4.7GB per disc in the next year) by a factor of twenty, first to 10GB per disc, ultimately up to 100GB per disc. In November 1997, the joint venture began a four year, \$21 million development program with approximately half of its funding to be provided by the ATP. The project combine four evolving technologies: Calimetrics' multilevel technology, which uses custom integrated circuits and specially written disks to increase data capacity and transfer rate; ECD's phase-change optical recording materials, which are used in rewritable DVD media; Polaroid's optics technology, which increases storage density per unit area and signal-to-noise (SNR) ratio; and Polaroid's versatile low-cost media processing approach. The technical challenges include the development of new coding and signal processing techniques to maximize disk capacity and performance; engineering of the media microstructure to support adequate SNR; and the use of special materials and processes to keep manufacturing costs down and control disk properties such as rewritability. The University of Arizona and Georgia Institute of Technology will provide support in optics, materials science, information coding, and signal processing.

This storage approach will compete directly with the projected roadmaps of conventional storage technologies (DVD-ROM/RAM, ASMO, etc.), potentially leveraging the progress in their component technologies to a high degree.

5.6. Probe storage

Probe storage techniques can generally be broken down into those that use a Scanning Tunneling Microscope (STM), Field Emission Probe (FEP), an Atomic Force Microscope (AFM), and Near

Field Scanning Optical Microscope (NSOM). These techniques can employ a variety of modulation techniques including topographic (mechanical), charge, magnetic, conductivity and optical modulation. They have been used and demonstrated with a large variety of materials including organic, ferro-electric, magneto-optic, magnetic and phase change media.

The Scanning Tunneling Microscope⁸⁴ takes advantage of the tunneling current that occurs across the gap between two conductors when they are held very close together (on the order of angstroms). With current flow, atoms may be deposited onto a substrate. With an STM it is actually possible to manipulate a single atom making the approach the highest density storage device demonstrated to date. However, most engineering efforts today concentrate on larger bit sizes (i.e. 10-20 nm size bits, giving a data density of over 1 Tbit/in²).

The main drawback to the STM is its slow data rate. Although individual bits may be recorded and read very quickly (using a voltage pulse of less than a nanosecond), the scanning of the STM probe is very slow. The height of the probe tip must be held very constant or else it will lose the tunneling current or crash into the substrate. Therefore scanning speeds are limited by the servo actuating speeds. Data rates as fast as 100 Kb/sec have been demonstrated.⁸⁵ Materials have been developed for read/write/erase applications using STM which undergo phase change (which changes their electrical conductivity) however these materials may suffer from fatigue and slow recording times.⁸⁶ The Atomic Force Microscope employs a cantilevered tip that is scanned across a sample. In this case the tip is usually in contact with the sample surface. Features on the surface of the sample cause deflections of the tip that can be detected optically by monitoring the back surface of the tip with a laser and split photodiode detector.

WORM writing with an AFM may be achieved using a thermo-mechanical process. The position of the AFM tip is placed concurrent with the focal point of a laser. When a pulse is sent out from the laser it locally heats the area around the tip causing the tip to sink into the polymer substrate creating a pit. Using this technique, pits as small as 100 nm have been recorded giving a recording density of roughly 30 Gbit/in². The recording speed for that experiment was 200 kb/sec and readout was up to 1.25 Mb/sec.

It may also be possible to produce inexpensive ROM replicas readable by an AFM. Disc "masters" have been fabricated using e-beam lithography.⁸⁷ The masters have then been used to replicate data in a photo-polymer. Features as small as 50 nm have been successfully replicated. Read/Write/Erase AFM storage has been done in a number of different materials including charge trapping materials such as nitride-oxide-semiconductor (NOS) layers⁸⁸, and also phase change materials the same as are using in optical discs.⁸⁹ The smallest bits that have been recorded are on the order of 75 nm giving a recording density of over 100 Gbit/in². Readout rates of 30 kb/sec have been achieved but predicted rates of up to 10 Mbit/sec may be possible. The major obstacle for AFM storage is tip wear. The tip is slowly worn down as it drags across the sample reducing the resolution of the readout.

Near-field optical recording includes any technology that allows one to surpass the areal density limit imposed by the diffraction of light in optical data storage. Near Field Scanning Optical Microscopy (NSOM) combats the effects of diffraction on the optical spot size by placing the pick up heads very near (about 50 nm above) the media. A near field probe can produce spots as small as 40 nm in diameter and conceptually can achieve areal densities in the order of 100 Gb/in². The problem is that the probe must be near contact with the medium making it difficult to prevent head crashes and support removable media.

One of the original near-field recording techniques was through the use of a tapered fiber.⁹¹ The tip of a fiber, which is smaller than the wavelength of the recording light, is positioned within 10 nm of the sample. Using magneto-optic materials, bits as small as 60 nm have been recorded and readout. The tapered fiber approach suffers from very low optical efficiency. STM, AFM and NSOM can all provide ultra high areal densities. However, the total area available for storage restricts the capacity per device (chip). Assuming a scan area of 1cm², 30nm effective spot size yields a total capacity of 10GB only. Yet using parallel access and MEMS technology, relatively fast access can be achieved. Considering the low power nature of these approaches, this type of probe storage is expected to have a potential market for applications that require portability.

In Japan, there has been significant progress in various areas of probe storage. This includes




Figure 37. Probe storage at Matsushita: Record breaking erasable spots of 10 nm diameter have been recorded in phase change media using STM.⁹⁰

hybrid STM/AFM systems, thermo-mechanical recording, automated Langmuir-Blodgett film deposition systems, new organic polymers, record breaking small recorded spots (Figure 37), and probe fabrication techniques that can lead to low cost manufacturable systems. Dr. Ohta at Matsushita, perhaps best underlined the expectations of Japanese researchers in probe storage by stating that "Phase Change media welcomes probe storage".

Research work in the U.S is equally active with IBM's efforts in atom herding and thermo-mechanical recording depicted in Figure 38. Thermo-mechanical recording at IBM using rotating plastic disk using an atomic force microscope tip. At selected locations, electrical current was




Figure 38. Thermo-mechanical recording at IBM using rotating plastic disk using an atomic force microscope tip. At selected locations, electrical current was pulsed through the tip causing the tip to be heated. The heat from the tip softens the substrate and the pressure from the tip.

pulsed through the tip causing the tip to be heated. The heat from the tip softens the substrate and the pressure from the tip.^{92,93} The majority of the work in the U.S. emphasizes demonstration of parallel access probe storage techniques using field emission probe techniques (e.g., work carried out at Hewlett Packard) as well as MEMS techniques that utilize phase change, ferro-electric and magnetic media (e.g., work carried out at Cornell^{94,95} and Carnegie-Mellon).

Probe storage may potentially offer compact, low power, moderate capacity storage for applications requiring portability (digital cameras, palm computers and digital assistants, portable information appliances). In this respect it will be competing as a replacement for FLASH memory, and the recent micro (1" diameter) removable magnetic hard drives.

6. 2-Photon Optical Storage Projections

6.1. Overview and progress

At Call/Recall, Inc., we are exploring two-photon materials and systems as a means to provide monolithic production of hundreds of data layers within inexpensive plastic disk media. Two-photon photochromic disk and cube media and parallel access reading and writing systems are being developed, with the goal of simultaneously providing more than 100 GB capacities and data transfer rates exceeding 1 Gb/s. This approach relies on recording bits in a volume by using two-photon absorption. Marks are written in the volume of an organic polymer disk only at points of temporal and spatial intersection of two beams with sufficient photon energies, one carrying information and the other specifying location, as in Figure 39.

Photochromic-dye-doped polymer example: spiropenzopyran-doped PMMA




Figure 39. Principle of 2-photon data recording and readout.

While single bits may be stored and recalled, parallel access of lines or planes of data is naturally accommodated in two-photon 3-D memories as in Figure 40, to provide increased data transfer rates. The recorded bits are read by fluorescence when excited by single-photon green illumination absorbed within the written spot volume. The results indicate no crosstalk between layers and excellent stability of the written bits at room temperature. This approach promises low-cost, high-density ROM or WORM disk media with hundreds of layers, and also low-cost 3D compact disk readers employing semiconductor green (or potentially red) lasers. In addition to the extremely large effective areal densities achievable with this monolithic 3-D disk approach, parallel access of the stored data can be naturally accommodated in this 3-D architecture, potentially leading to a significant increase in data transfer rate.




Figure 40. Orthogonally addressed 2-photon recording architectures.

Our recent media experiments have concentrated on optimizing the 2-photon photochemistry of our stable ROM/WORM materials, though we are also exploring erasable photochemistries.⁹⁶ Manufacturing and environmental testing issues are also being investigated, and thick (5-10 mm) disk injection molding and temperature/humidity stress studies are underway. Most of our system experiments have used our ROM/WORM material, in the form of 10x10x10 mm³ cube samples, as well as in 1 cm x 3 cm diameter disks. These are fabricated by incorporating a desired concentration of active molecules in liquid MMA, polymerizing the solution, followed by cutting and polishing. Masks are illuminated by 1064 nm, 35 psec pulses from a mode-locked Nd:YAG laser and imaged into the media to form the data planes. The frequency doubled green beam from the laser (532 nm) is incident orthogonally to the image path and anamorphically focused to a sheet which is spatially and temporally aligned with the IR image plane within the media. Retrieval and analysis of the data is performed with approximately 500 μ W of CW illumination at 543 nm from a HeNe laser introduced along the green path of the recorder. The induced fluorescence is imaged onto a CCD camera and analyzed with custom software to perform the bit thresholding decisions and calculate average signal separations and spatial bit-error-rates. Retrieval of images from the memory has also been demonstrated in a portable ROM system using a simple stepper motor driven stage, a 200 μ W HeNe laser, and a low cost video camera (see).⁹⁷

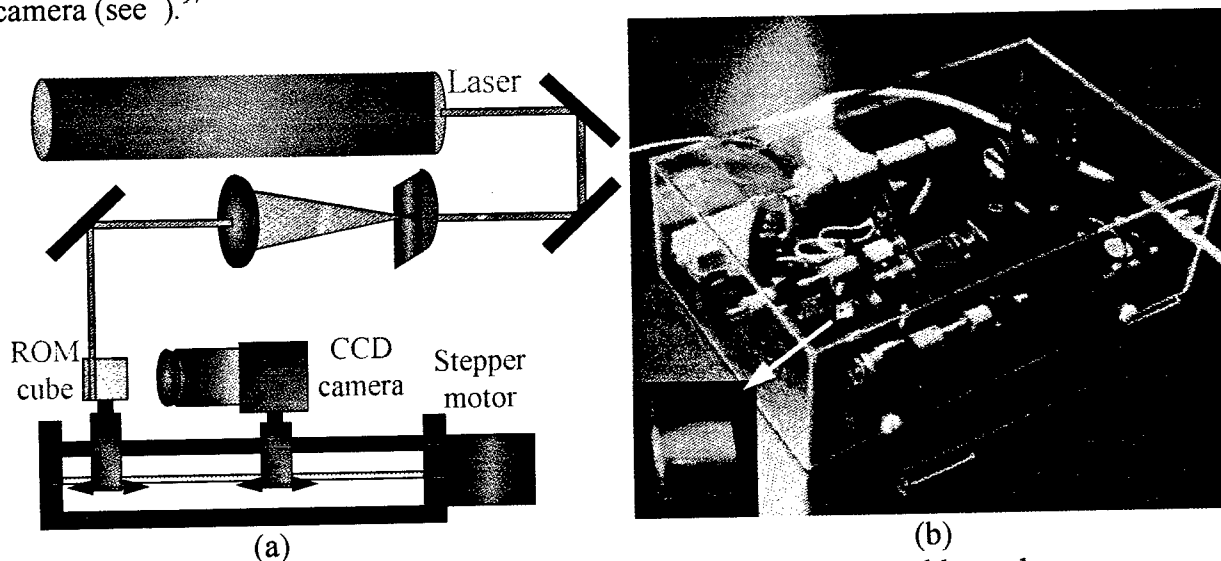


Figure 41. Portable ROM reader: (a) schematic, (b) portable reader.

In 1996, 100 data planes having 10^4 random bits/plane were recorded in the ROM cube media.⁹⁸ The planes were recorded on an 80 μm pitch, as shown in Figure 42. An autogain algorithm was implemented to remove the effects of the laser mode non-uniformity. This resulted in 493 errors out of 960,400 measured bits, mostly due to small scratches and dust particles on the cube surface. The average spatial SNR for the 100 planes was 8.08:1, with local SNRs greater than 19:1 measured within individual data planes, indicating the potential impact of sample fabrication improvements. The SNR of 19:1 corresponds to a spatial BER of less than 9×10^{-13} .




Figure 42. Example image, data plane, and side view of the 100 layer, 8mm thick cube.

Recent work has also demonstrated storage and recall of information from rotating disk media.⁹⁹ We are currently exploring several disk-based reading and writing architectures, and we have constructed the automated disk publishing setup shown in Figure 42

This setup uses the same laser as is used to record cube samples, but provides 0.1 μm positioning accuracy of the disk, and allows several disk recording and readout geometries to be explored. As many as 17 layers have been recorded in a 5 mm thick, 25 mm diameter disk using this setup. To examine dynamic reading effects from a spinning 3-D disk, we have constructed the disk test stand shown in Figure 43. This simple disk test stand uses a diode pumped solid state laser (doubled to provide 532 nm green illumination), aspheric glass CD objective lenses, and a DC servo motor, and can read disks in either a colinear (illumination & data through same lens) or orthogonal geometry (illumination & data through different lenses).




Figure 43. (a) Automated disk recorder, and (b) disk test stand.

Figure 44 shows the test stand, and a side view of the disk with all of the layer pairs illuminated. Preliminary testing has demonstrated readout of data tracks at 1200 rpm, to investigate interplane crosstalk, pairs of layers were recorded in a disk with separations of 20 to 200 μm between pairs.

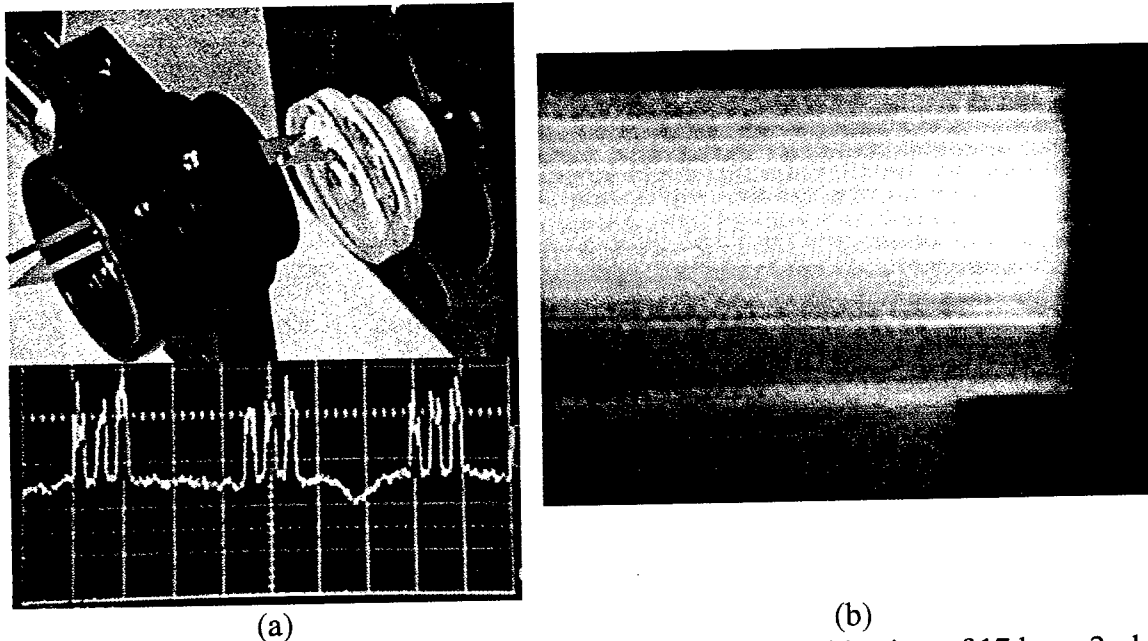


Figure 44. (a) Colinearly illuminated disk and data signal, and (b) side view of 17 layer 2-photon 3D disk.

6.2. Projections

6.2.1. Recording

Recording rate projections are largely dependent on the increasing availability, and decreasing cost of short pulse lasers. A variety of laser systems in this category is now available, and simple theoretical models predict that some of these lasers should provide sufficient recording rates to be competitive from a performance standpoint. For cost competitiveness, the lasers system costs must also continue to decrease, by more than one order of magnitude.




Figure 45. Preliminary recording results achieved using a 76 MHz, 200 fs pulse length near infrared laser.

Preliminary experiments using high repetition rate lasers (e.g.- Coherent's MIRA or Vitesse lasers sources) show promise for high rate recording of low crosstalk multilayers, as shown in Figure 45. Other lasers, such as the Lightwave Electronic LWE 131¹⁰⁰, the Energy Compression Research PML400¹⁰¹, the Microlase DPM-PC¹⁰², the sfim ODS TinyGreen¹⁰³, or the TimeBandwidth Products Tiger-200 or Picolo-500¹⁰⁴, also offer potential solutions for high repetition rate recording sources. However, these sources also need to be reduced in size and cost by 1 or 2 orders or magnitude, while maintaining or improving their performance. Given the rapid rate of progress in these laser systems (most of these companies did not exist 5 years ago), such improvements are not unlikely.

6.2.2. Reading: the FROST architecture

Advanced 3D reading technologies and systems are currently being developed under a 4 year




Figure 46. Fast Readout Optical Storage Technology (FROST) approach

DARPA sponsored program which funds a research consortium consisting of Call/Recall, Hewlett-Packard, Irvine sensors, and the University of Southern California via a Technology Investment Agreement administered by the Air Force Research Labs. The Fast Readout Optical Storage Technology (FROST) program seeks to combine the inherent advantages of multilayer optical storage media in terms of density and removability with VLSI Photonic techniques to develop high-throughput parallel optical readout systems which will enable orders of magnitude improvement over conventional storage systems in future military and commercial digital systems. The end-goal of the program is to demonstrate a bit-oriented, volumetric removable optical disk system with a parallel readout head providing both high capacity and high-throughput. The techniques employed will have the potential for more than 100 GB per disk capacities and greater than 10 Gb/s transfer rates, though the goal of first prototype is a 50 GByte disk system with 2 Gb/s throughput. Emphasis will be placed on the development of the component technologies needed for the various subsystems and on the integration of these novel components into several prototype demonstrations. To achieve this goal, material processing and device design developments will cover photochromic polymers, vertical cavity semiconductor diode laser arrays, optical illumination and 3-D imaging systems, micro-mechanical optical actuator chips, high-performance detector arrays, and parallel data channel electronics. In addition, special attention will be placed on the packaging techniques required to assemble all the components into a cost-efficient means of retrieving optical data stored in a 3-D format.

In the FROST system shown in Figure 46, the format of the data stored within the media can be viewed in two ways: as a large stack of layers, similar to conventional 2-D disk layers, or as a collection of data planes, tilted with respect to the flat disk surfaces. The tilted plane view is better matched to the actual readout scheme. In this view, each plane is rectangular, longer in the disk thickness direction than radially. Thus in the radial direction, what were the “tracks” in serial 2-D disks, are grouped together to form “supertracks”, which extend into the depth of the disk to form the data plane. The VCSEL beams in the illumination array are combined and focused into a thin sheet to edge-illuminate a tilted data plane within the disk, thus causing the bits to fluoresce. As the disk rotates, different data planes are brought into coincidence with the fixed sheet of light. The fluorescent radiation from the recorded bits in the data plane is collected by the depth transfer optics, and imaged onto the data sensor array. Tracking and focus control is achieved by monitoring the detected signals of the corners of the array, and the fine focus/tracking signals are generated on the head and sent directly to micromachined actuators within the optical path of the Depth Transfer Objective (MAMA FACT actuators). The detected data are amplified, equalized, and thresholded, and error control methods, address and synchronization interpretation, channel multiplexing, and data filtering is performed. Stacked neo-chip integration and potting techniques may be used to package the electronics, optoelectronics, and optics in to modular subassemblies, and precision alignment structures fabricated on a silicon microbench will enable passive alignment and assembly of these modules.

6.2.3. FROST performance and scaling projections

Table 4 summarizes the near and long term performance potentials of the FROST approach. The 3-dimensional media density and the overall media volume determine single disk capacity. The practical 3-D density for this demonstration can be estimated from the optical recording and readout resolutions and the head positioning accuracy. The optical resolution is determined by the wavelength of light and the numerical aperture (NA) of the optics to be used. Since the proposed 2-photon media will use 1064nm light for the data image, versus 650-700nm light for the readout, the recording resolution can be used as a worst-case. If we assume a recording system NA of 0.6 (equivalent to the DVD readout optics, the DVD recording optics and wavelengths are of significantly higher resolution), we achieve a lateral Airy disk radius of 1 μm for the data marks. Since low cost CD tracking and focus servo systems routinely position their heads over marks smaller than this, we can assume that head positioning will not be the limiting factor. While in principle, the data plane spacing could be similar to the lateral mark spacing (especially if the proposed 3-D equalization techniques are employed), for the purpose of this estimate we will assume a 10 mm data plane pitch to avoid recorded interplane crosstalk effects. Assuming DVD encoding efficiencies, a conservative first order estimate of the bit volume is 5 $\mu\text{m}^3/\text{bit}$. Again assuming a conservative DVD disk area utilization, a 120mm diameter by 10mm thick 3DCD disk provides a raw capacity of 1.8 terabits. Accounting for encoding, ECC, and synchronization/addressing overhead results in a user capacity of just over 100 GB/disk. Similarly, using the above disk dimensions, and assuming the proposed 32x128 detector sensor array, the rotation rate required to provide 4 Gb/s aggregate raw data transfer rate is only 1300

rpm. Again assuming DVD encoding overheads, the user achievable user data rate at this rotation speed can be greater than 2 Gb/s.

Several paths to scaling this performance in the future are clear. First, the recording wavelength may be shortened. Development of media sensitive to blue wavelengths is planned, which, due to the long wavelength recording of the current media, would result in a 19x capacity increase. Second, similar to DVD publishing, the publishing facility for 3DCD ROMs can use more expensive, higher NA optics, e.g.- if NA=0.9 optics were employed, a 3.4x volumetric capacity improvement results. Together, these result in a 64x increase in capacity, or a 6.4 TB user capacity disk. The 1-dimensional data density also acts to scale data transfer rate, as do the rotation speed and sensor parallelism. Since the bit's linear size and the rotation rate limit the sensor integration time, this also necessitates increased sensor performance and/or greater fluorescence signal strength from each bit. Assuming only a 2x increase in sensor performance, a 4x increase in fluorescence strength, a 2x in the rotation rate, and increasing the sensor array parallelism from 32x128 to 128x512 (still much smaller than the 1k x 1k APS sensor arrays projected within the next 3 years), results in sustained user data transfer rates exceeding 65Gb/s. Finally, given the high reliability and low projected system and media costs of FROST 3DCD systems, using them in bulk disk arrays appears to be a highly effective means to providing capacities in the 100 TB-1PB regime, with associated data transfer rates in the 1-10Tb/s range.

Continued progress along this roadmap in a cost-effective manner will, of course, require the continued evolution of the system components (media, lasers, optics, sensors, and parallel electronics), as discussed in section 7.1.

Table 4. 3DCD FROST system performance scaling projections

	Proposed	Projected	Factor
Recording wavelength (nm)	1064	400	2.7x
Recording NA	0.6	0.9	1.5x
Mark volume (μm^3)	2x2x10	0.5 x 0.5 x 2.5	64x
bits/mark (in 3D)	2x2x2	2x2x2	1x
Disk dimensions (mm^3)	120 diam. x 10 thick	120 diam x 10 thick	1x
Raw capacity per disk (Tbits)	1.8	115	64x
Encoding efficiency rate	0.5	0.5	1x
User capacity per disk (GB)	100	6,400	64x
Rotation rate (rpm)	1300	2600	2x
Nom. signal strength (nW)	5	20	4x
Nom. sensor sensitivity (nW)	5	2.5	2x
Sensor parallelism	32x128	128x512	16x
Raw data transfer rate (Gb/s)	4	512	256x
User data transfer rate (Gb/s)	2	256	128x




Figure 47. Target areal densities for mass-market advanced optical storage technologies.
 (Source: Marvin Keshner, HP Labs, in the High Density Optical Data Storage Workshop,
 San Diego, CA, March 2, 1998.)

7. Windows of opportunity

As can be seen from Figure 47, it will be increasingly difficult to exceed 100Gb/in² or 100GB/disk (12cm diameter) with conventional techniques. Based on the roadmaps presented in earlier sections, a window of opportunity for alternative storage technologies should open by 2005. Thus, it is important to study the potential performance characteristics of new emerging technologies and understand their application potentials.

The areal density of conventional data storage systems will encounter certain physical or engineering limits within the next decade. For example, the diffraction limit of light puts a limit on the practical spot size in optical disc systems; the thermo-magnetic limit sets up a maximum areal density for longitudinal magnetic recording used in present-day hard drives. In view of these limits it is expected that new, more unconventional technological solutions may become necessary. These solutions may entail the further increase of the areal density by using, for example, near field optics or perpendicular magnetic recording. The solutions may also take advantage of additional available degrees of freedom such as those provided by volumetric storage techniques. For example, holographic and two-photon optical storage approaches enable data to be recorded in the volume of a medium rather than just its surface, thereby achieving much greater volumetric densities and data locality. The solutions may also be derived from significantly different technology platforms that may lead to new types of data storage systems. For example, Micro Electro Mechanical Systems (MEMS) technology can enable probe storage techniques to achieve 100X better areal densities than projected for conventional technologies. In a more distant future biologically inspired nano-structures may be envisioned to play important roles in storage systems as well. Finally, solutions may result from indirect impact of new technology platforms. For example, the advent of low cost vertical cavity surface emitting laser arrays and of micro-mechanical actuators may significantly affect the cost and design of pick-up heads enabling fast access parallel recording and readout. Several emerging technologies are especially well placed to capture new markets from conventional storage technologies by the second half of the next decade. Some potential technology/application matches are enumerated in Table 5.

Table 5: Emerging technologies and applications

Technology	Applications	Performance
Structured magnetic media SIL lens near-field MO disks	General purpose, computing	200GB, ms access
Volumetric disks	Advanced video, removable	200GB-1TB, 1Gb/s data rate
Thermo-mechanical probe (disks)	Mastering	1TB and more, Mb/s data rate
Probe storage (chips)	Portable video applications	10GB, μ s access, 100Mb/s

7.1. Windows of opportunities for Two-Photon Disk Storage

The application pull that is expected to continue on one hand and the limits of conventional storage technologies being in sight on the other, opens windows of opportunities for alternative storage technologies such as two-photon disk storage. Indeed the importance of storing information in the volume of a storage medium rather than only on its surface has become evident over the last few years. Perhaps, the power of two-photon recording technology is best illustrated in Figure 48 and Figure 49, which indicate that by increasing the number of storage layers, two-photon recording will be able to exceed areal densities beyond 100Gb/in² and certainly approaching 1Tb/in².

At this point in time two-photon recording appears to be especially well positioned to carry-on DVD like products to new performance levels because of its inherently low cost media enabling many storage layers. Indeed, at the present time, this technology can be used to offer products with a common format for ROM as well as WORM applications. The main roadblock that may delay the entry point of two-photon recording based WORM devices appears to be the recording lasers, which at this time are too costly.

An important development direction for two-photon WORM type of storage systems appear to be the development of necessary techniques to boost the data transfer rates above 1Gb/s. To this end the investigation of parallel access techniques seems critical. To this end, the development of optical sensor arrays will be critical for most parallel access storage systems. MEMS actuators may provide a low cost solution for certain types of aberration correction and tracking and focusing functions. Parallel ECC & data filter processing will need to be lowered in cost.

Another important direction for research in two-photon storage systems is the development of erasable materials. If erasable two-photon media could be developed over the coming few years,




Figure 48. Possible road map for alternative technologies

(Source: Derived from various data obtained from TeraStor Web site, OITDA storage roadmap, private communications with Holoplex researchers and Call/Recall Inc. internal reports)




Figure 49. Potential evolution of capacity vs. transfer rate and comparison with conventional techniques (Based Call/Recall 1997 storage survey)

than this technology will be able to address the very important random access read/write applications and provide an alternate solution as a low cost hard disk replacement.

7.1.1. Potential product evolution path

To launch a series of new 3DCD products by the second half of the next decade (e.g., as DVD replacements), there are immediate corporate R&D needs in two areas: i.) accessing some of the required components to develop next generation prototypes ii.) attracting new partners from the content providers community and media manufacturers. Additionally, this approach may also lead to new products in components and sub-systems market segments for optoelectronics in general. For example, optical sensor arrays, MEMS actuators, parallel ECC & data filter processing stacks, and visible VCSEL arrays may find applications in various optoelectronics market segments other than data storage.

Commercial applications will facilitate the cost-effective transitioning of these technologies into military markets. For example, the application of VCSELs to the next generation optical storage systems also offers new opportunities in systems exploiting parallel access. The inherent redundancy of VCSEL arrays used for illumination can provide for higher reliability necessary for military applications. More robust tracking and focus will be made possible by the availability of extra emitters dedicated to specific tasks. Finally, in detector arrays, consumer application cameras will drive the array sizes and basic functionality, but military and storage applications will stress additional characteristics, including sensitivity (for low light imagery and tracking), speed (for autonomous navigation and UAV), local-pixel processing (for novelty filtering, and data reduction and fusion).

8. Appendix I: Projection Validation Workshop

The High Density Optical Data Storage Workshop, took place at the Wyndham Hotel (Sorrento Valley) in San Diego, March 2nd, 1998. The objective of this meeting was to bring together invited participants from government, industry, and academia to provide input on how optical storage technologies can exceed 60Gb/in² data densities and which technologies are worth pursuing in this goal. Special emphasis was placed on emerging applications as well as on non-conventional approaches. This workshop was also scheduled just prior to a study sponsored by WTEC on Optical Data Storage in Japan in order to facilitate interactions between the attendees of this workshop and the members of the WTEC panel. The one day workshop also provided an opportunity for the initial results of this program to be presented to external experts and for their input to be incorporated into the market trends, conventional technology, and advanced technology portions of this report.

8.1. Workshop agenda

WORKSHOP ON HIGH DENSITY OPTICAL DATA STORAGE

Agenda

8:30am to 12pm

Welcome and Introduction

Sadik Esener (Call/Recall, Inc.)

Emerging Military Applications

Fred Haritatos (Air Force Research Labs - Rome)

Applications and Signal Processing for Optical Data Storage

C. Kuznia and A. Sawchuk (USC)

break

Present status and potentials of two-photon storage

Frederick McCormick (Call/Recall, Inc)

Present status and potentials of PSHB storage

Eric Maniloff (Templex)

Present status and potentials of holographic storage (Canceled)

Marc Neifeld (Univ. of Arizona)

Summary of holographic and other emerging storage technologies

Sadik Esnser (Call/Recall, Inc.)

Summary of Current Ideas in HDD

Marvin Keshner (HP)

Open Commentary/Discussion session (5-10 min per speaker)

12-1:30pm Lunch and explanations of problem specifications (see below)

1:30-4:00pm Working groups

4-4:20pm break

4:20-5:00pm Summary of working groups

Problem Statement:

Future applications (beyond 2005) require affordable data storage systems capable of providing high capacity (in excess of 60 Gb/in² effective areal densities) and high performance. As an example set of system requirements, consider a system providing: 250 GB erasable random access capacity, with < 20 ms access time, 10 Gb/s data transfer rate, inexpensive removable media, and robust (Mil-Spec environment) operation within a 15 x 15 x 5 cm³ volume. Storage systems developed to meet these application requirements will need to integrate advanced components including media, lasers, data sensors, actuators, microoptics, and parallel optoelectronic signal processing, into cost-effective, robust packages. During the workshop, working groups will determine which components are most critical (e.g.- rank order) to meet these requirements, as well as to positively impact US economic competitiveness in the optical storage arena.

8.2. Workshop attendees

Name	Affiliation	Email Address
Al Jamberdion	Rome Labs	alj@rl.af.mil
Alexander Dvornikov	Call/Recall, Inc.	advornik@call-recall.com
Andrew Sheckl	University of Cincinnati	
Barbara Capron	Boeing	
Bernie Clark	Rome Labs	clarkeb@rl.af.mil
Charlie Kuznia	USC	kuznia@sipi.usc.edu
Elias Towe	DAPRA	etowe@darpa.mil
Emily Wesseling	JPL	
Eric Maniloff	Templex Technologies	
Fai Mok	Holoplex	fai@holoplex.com
Fred Haritatos	Rome Labs	haritatosf@rl.af.mil
Gary Sunata	Rome Labs	
Glenn Sincerbox		
Ilkan Cokgor	HP Labs	
Marvin Keshner	HP Labs	
Pericles Mitkas		
Pete Costianes	Rome Labs	
Peter Rentzepis	Call/Recall, Inc.	pmrentze@call-recall.com
Philippe Marchand	Parallel Solutions	pmarchand@ucsd.edu
Rick McCormick	Call/Recall, Inc.	fmac@call-recall.com
Sadik Esener	Call/Recall, Inc.	sadik@call-recall.com
Sandy Sawchuk	USC	sawchuk@sipi.usc.edu
Shay Capehart	Rome Labs	
Shaya Feinman	UCSD	fainman@ece.ucsd.edu
Suzanne Wakelin	Optivision, Inc.	
Volkan Ozguz	Irvine Sensors	vozguz@irvine-sensors.com
W.E. Moerner	UCSD	

9. References

- ¹ <http://www.datathink.com/menu.html>
- ² <http://www.internationaldatacorp.com>
- ³ <http://www.cinesite.com>
- ⁴ For additional information on data warehousing and mining, consult the URL <http://pwp.starnetinc.com/larryg/techeval.html>
- ⁵ Internet Week 4, Nov. 1997.
- ⁶ Bob Violino, Informationweek Daily, Jan. 25, 1999, <http://www.informationweek.com>
- ⁷ Martin J. Garvey, Informationweek Daily, Jan. 25, 1999, <http://www.informationweek.com>
- ⁸ Hinteregger, H.F., Rogers, A.E.E., Capallo, R. J., Webber, J.C., Petrachenko, W.T., and Allen, H. 1991. A High Data Rate Recorder for Astronomy. IEEE Transactions on Magnetics. Volume 27, Number 3, May 1991.
- ⁹ <http://web.fie.com/htdoc/fed/afr/sab/any/text/any/uavstud.htm>. For overview information on UAVs, see for example: <http://www.erols.com/auvsicc/uav.htm> or <http://www.hugh.demon.co.uk/Book22.html>
- ¹⁰ <http://web.fie.com/htdoc/fed/afr/sab/any/text/any/afrttech.htm>
- ¹¹ Wing Commander Mark Lax and Wing Commander Barry Sutherland, "An Extended Role for Unmanned Aerial Vehicles in the Royal Australian Air Force," Air Power Studies Centre paper number 46, <http://www.defence.gov.au/apsc/publish/paper46.htm>
- ¹² Jane's Defense Weekly, 5 August 1995, p. 32.
- ¹³ An ordinary 203 mm x 203 mm picture has 1024x1024 pixels with each pixel constituting 16 bits, thereby giving a total of approximately 2.1 megabytes of data. Multispectral imagery at higher resolutions increase this number dramatically.
- ¹⁴ Jane's Defense Weekly, 12 August 1995, p. 38
- ¹⁵ Michael A. Dorheim, 'New Sensors Show Two Paths of Reconnaissance', Aviation Week & Space Technology, 10 July 1995, pp. 42-44
- ¹⁶ UAV Force Concept, Aviation Week & Space Technology, 10 July 1995, p. 43.
- ¹⁷ Jane's Defence Weekly, 2 September 1995, p. 3.
- ¹⁸ <http://www.dmsi.mil/docslib/mspolicy/msmp/1095msmp/Chap2.html#2A>
- ¹⁹ See, for example, <http://www.stricom.army.mil/PRODUCTS/WARSIM/home.html>, and <http://www.jsims.mil/home/>
- ²⁰ Holodeck is a term extracted from the TV Star Trek series referring to a computer generated virtual reality room which allows one to experience the sense of being in any environment, limited only by their imagination.
- ²¹ Lt Col Robert L. Atkins, Jr., Lt Col Duane A. Lamb, Ms Marlene L. Barger, Maj Larry N. Adair, Maj Michael J. Tiernan, "2025 In-Time Information Integration System (I3S)," Air Force 2025, April 1996, <http://www.au.af.mil/au/2025/>
- ²² James Kitfield, "Trading Bullets for Bytes," Government Executive, June 1994, 18.
- ²³ Brenda Laurel, "Virtual Reality," Scientific American, September 1995, 70.
- ²⁴ Lt Col Gregory J. Miller, Maj Kurt C. Fecht, Maj Barbara Jeffs, Maj Laura J. Sampsel, Maj Patrick J. Smith, "Virtual Integrated Planning And Execution Resource System (VIPERS): The High Ground Of 2025," August 1996, <http://www.au.af.mil/au/2025/>
- ²⁵ Rome Laboratory FY95 Command Control Communications and Intelligence C3I Technology Area Plan (Wright-Patterson Air Force Base, Ohio, Headquarters Air Force Materiel Command, June 1994).
- ²⁶ Lt Col David M. Snyder, Maj Penny J. Dieryck, Maj Wesley W. Long, Maj Thomas G. Philipkosky, Lt Cmdr Ronald Reis, "Joint Readiness Assessment and Planning Integrated Decision System (JRAPIDS): Combat Readiness and Joint Force Management for 2025," Air Force 2025, <http://www.au.af.mil/au/2025/>
- ²⁷ Barbara A Capron, "Optical Memory Application Assessment," Technical Final Report, CDRL CDLIN A002, Contract 30602-95-C-0270.
- ²⁸ The prime contractor for ATIMS is GDE systems, <http://gdesystems.com/MPS/SlipSheets/ATIMS.html>

- ²⁹ New World Vistas: Air and Space Power for the 21st Century, <http://www.fas.org/spp/military/docops/usaf/vistas/vistas.htm>
- ³⁰ <http://gcmd.gsfc.nasa.gov/eos/>
- ³¹ <http://www-physics.ucsd.edu/~chacker/Smart/SMARTmtgnotes.html>
- ³² Ernest Raia, Data Storage magazine, pp.5, December, 1998.
- ³³ Investor's Business Daily 2 Dec 1998.
- ³⁴ Lionel C. Kimerling, "The economics of science: from photons to products," Optics and Photonics News, Optical Society of America, pp. 19, October, 1998.
- ³⁵ Janos Kovacs, "Read-channel technology for the >200-Mbits/ era", Computer Design, pp. 108, Dec 1996.
- ³⁶ The National Technology Roadmap for Semiconductors, Semiconductor Industry Association, p. B2 (1994).
- ³⁷ Regan Zane and H. Scott Hinton, "Progress Report on Optical Interfaces for Advanced Storage Systems," Photonic Systems Group, University of Colorado, September 4, ¹⁹⁹⁶.
- ³⁸ Roadmaps for Parallel Optoelectronic Interconnects, Optoelectronics Industry Development Association (1994).
- ³⁹ <http://www.zenresearch.com/index.html>
- ⁴⁰ Goto, Kenya, "High bit rate and Tera bytes optical memory in a disk system," In Optical Data Storage '97, Henryk Birecki, James Z. Kwiecien, Editors, Proceedings of SPIE Vol. 3109, 192-95 (1997).
- ⁴¹ <http://www.quinta.com/>
- ⁴² B. D. Terris, H. J. Mamin, D. Rugar, W. R. Studenmund, and G. S. Kino, "Near-field optical data storage using a solid immersion lens," Appl. Phys. Lett. **65**, 388 (1994).
- ⁴³ B. D. Terris, H. J. Mamin, and D. Rugar, "Near-field optical data storage," Appl. Phys. Lett. **68**, 141 (1996).
- ⁴⁴ I. Ichimura, S. Hayashi, and G. S. Kino, "High-density optical recording using a solid immersion lens," Appl. Opt. **36**, p. 4339 (1997).
- ⁴⁵ M. Mansuripur, "Tutorial on near field recording," presented at MORIS '99, Monterey, CA, Jan. 11, 1999.
- ⁴⁶ <http://www.holoplex.com>
- ⁴⁷ D.L. Staebler and J.J. Amodei, "Coupled Wave Analysis of Holographic Storage in LiNbO₃," Journal of Applied Physics, **34**(3):1042-1049, 1972.
- ⁴⁸ K. Curtis, A. Pu, and D. Psaltis, "Method for Holographic Storage using Peristrophic Multiplexing," Optics Letters, **19**(3):993-994, 1994.
- ⁴⁹ D. Psaltis, et al., "Holographic Storage using Shift Multiplexing," Optics Letters, **20**(7):782-784, 1995.
- ⁵⁰ A.C. Strasser, et al., "Procedure for Recording Multiple-Exposure Holograms with Equal Diffraction Efficiency in Photorefractive Media," Optics Letters, **14**(1):6-8, 1989.
- ⁵¹ J.J. Drolet, et al., "Compact, Integrated Dynamic Holographic Memory with Refreshed Holograms," Optics Letters, **22**(8):552-554, 1997.
- ⁵² See for example: Y. Qiao, S. Orlov, and D. Psaltis, "Electrical Fixing of Photorefractive Holograms in SBN," Optics Letters, **18**(2):1004-1006, 1993, or V. Leyva, A. Agranat, and A. Yariv, "Fixing of a Photorefractive Grating in KTN by Cooling Through the Ferroelectric Phase Transition," Optics Letters, **16**(8):554-556, 1991., or J.F. Heanue, et al., "Digital Holographic Storage Incorporating Thermal Fixing in Lithium Niobate," Optics Letters, **21**(19): 1615-1617, 1996.
- ⁵³ F.H. Mok, G.W. Burr, and D. Psaltis, "System Metric for Holographic Memory Systems," Optics Letters, **21**(12):896-898, 1996.
- ⁵⁴ D.A. Waldman and H.Y.S. Li, "Determination of Low Transverse Shrinkage in Slant Fringe Gratings of a Catatonic Ring-Opening Volume Hologram Recording Material," Proceedings of SPIE, **3010**:354-372, 1997.
- ⁵⁵ K. Curtis, "Digital Holographic Storage," invited presentation at the Optical Society of America Annual Meeting, Paper ThR1, Long Beach, CA, 1997.
- ⁵⁶ G.W. Burr, F.H. Mok, and D. Psaltis, "Large Scale Volume Holographic Storage in the Long Interaction Length Architecture," Proceedings of SPIE, **2297**:402-414, 1994.
- ⁵⁷ H.Y.S. Li and D. Psaltis, "Three-Dimensional Holographic Optical Disks," Applied Optics, **33**(17):3764-3774, 1994.
- ⁵⁸ A. Pu and D. Psaltis, "High Density Recording in Photopolymer-Based Holographic 3D Disks," Applied Optics, **35**(14):2389-2398, 1996.

- ⁵⁹ S. Campbell, X. Yi, and P. Yeh, "Hybrid Sparse-Wavelength Angle-Multiplexed Optical Data Storage System," *Optics Letters*, 19(24):2161-2163, 1994.
- ⁶⁰ J.F. Heanue, M.C. Bashaw, and L. Hesselink, "Sparse Selection of Reference Beams for Wavelength- and Angle-Multiplexed Volume Holography," *Journal of the Optical Society of America A*, 12(8):1671-1676, 1995.
- ⁶¹ J. Lembcke, C. Denz, and T. Tschudi, "General Formalism for Angular and Phase-Encoded Multiplexing in Holographic Image Storage," *Journal of Optical Materials*, 4:428-432, 1995.
- ⁶² C. Denz, et al., "Potentialities and Limitations of Hologram Multiplexing by using the Phase-Encoding Technique," *Applied Optics*, 31(26):5700-5705, 1992.
- ⁶³ M.C. Bashaw, et al., "Cross-Talk Considerations for Angular and Phase-Encoded Multiplexing in Volume Holography," *Journal of the Optical Society of America B*, 11(9):1820-1836, 1994.
- ⁶⁴ L. Hesselink and M.C. Bashaw, "Optical Memories Implemented with Photorefractive Media," *Journal of Optical and Quantum Electronics*, 25(9):611-661, 1993.
- ⁶⁵ A. Pu and D. Psaltis, "High Density Recording in Photopolymer-Based Holographic 3D Disks," *Applied Optics*, 35(14):2389-2398, 1996.
- ⁶⁶ G. Barbastathis, et al., "Shift Multiplexing with Spherical Reference Waves," *Applied Optics*, 35(14):2403-2417, 1996.
- ⁶⁷ K. Curtis, "Digital Holographic Storage," invited presentation at the Optical Society of America Annual Meeting, Paper ThR1, Long Beach, CA, 1997.
- ⁶⁸ H.Y.S. Li and D. Psaltis, "Three-Dimensional Holographic Optical Disks," *Applied Optics*, 33(17):3764-3774, 1994.
- ⁶⁹ <http://www.rsc.rockwell.com/html/seed.html>
- ⁷⁰ V.B. Markov, "Spatial-Angular Selectivity of 3D Speckle Wave Holograms and Information Storage," *Journal of Imaging Science and Technology*, 41(4):383-388, 1997.
- ⁷¹ See for example the Kodak C4 camera.
- ⁷² M.A. Neifeld and W.C. Chou, "Information Theoretic Limits to the Capacity of Volume Holographic Optical Memory," *Optics Letters*, 36(2):514-516, 1997.
- ⁷³ M.A. Neifeld and M. McDonald, "Lens Design Issues Impacting Page Access to Volume Optical Media," *Optics Communications*, 120(1995):8-14, 1995.
- ⁷⁴ H.Y.S. Li and D. Psaltis, "Three-Dimensional Holographic Optical Disks," *Applied Optics*, 33(17):3764-3774, 1994.
- ⁷⁵ D.A. Waldman and H.Y.S. Li, "Determination of Low Transverse Shrinkage in Slant Fringe Gratings of a Catatonic Ring-Opening Volume Hologram Recording Material," *Proceedings of SPIE*, 3010:354-372, 1997.
- ⁷⁶ M. Liphard, et al., "High Performance Photorefractive Polymers," *Science*, 263:367-369, 1994.
- ⁷⁷ B. Volodin, et al., "Highly Efficient Photorefractive Polymers for Dynamic Holography," *Optical Engineering*, 34(8), 2213-2223, 1995.
- ⁷⁸ Wenhai Liu, Ernest Chuang, Demetri Psaltis, "Read/Write Holographic Memory versus Silicon Storage," in *Advanced Optical Memories and Interfaces to Computer Storage*, Pericles A. Mitkas, Zameer U. Hasan, Editors, *Proceedings of SPIE Vol. 3468*, 2-7 (1998).
- ⁷⁹ <http://opticb.uoregon.edu/~mosswww/memory/shm.html>
- ⁸⁰ <http://www.templex.com/>
- ⁸¹ T. W. Mossberg, *Opt. Lett.* 7, 77 (1982); *Opt. Lett.* 17, 535 (1992).
- ⁸² See also: X. A. Shen and R. Kachru, *Opt. Lett.* 18, 1967 (1993) and M. Mitsunaga, R. Yano, and N. Uesugi, *Opt. Lett.* 16, 1890 (1991).
- ⁸³ <http://www.calimetrics.com/>
- ⁸⁴ I. D. M. Eigler and E. K. Schweizer, *Nature* 344, 524 (1990).
- ⁸⁵ H. J. Mamin, B. D. Terris, L. S. Fan, S. Hoen, R. C. Barrett, and D. Rugar, "High-density data storage using proximal probe techniques," *IBM J. Res. Develop.* 39, 681 (1995).
- ⁸⁶ A. Sato and Y. Tsukamoto, "Nanometre-scale recording and erasing with the scanning tunnelling microscope," *Nature* 363, 431 (1993).

- ⁸⁷ B. D. Terris, H. J. Mamin, M. E. Best, J. A. Logan, D. Rugar, and S. A. Rishton, "Nanoscale replication for scanning probe data storage," *Appl. Phys. Lett.* **69**, p. 4262 (1996).
- ⁸⁸ I. Fujiwara, S. Kojima, and J. Seto, "High Density Charge Storage Memory with Scanning Probe Microscopy," *Jpn. J. Appl. Phys.* **35**, 2764 (1996).
- ⁸⁹ H. Kado and T. Tohda, "Nanometer-scale recording on chalcogenide films with an atomic force microscope," *Appl. Phys. Lett.* **66**, 2961 (1995).
- ⁹⁰ H. Kado and T. Tohda, "Nanometer-scale recording on chalcogenide films with atomic force microscope", *Appl. Phys. Lett.* **66**(22), p.2961 May 1995
- ⁹¹ E. Betzig, J. K. Trautman, R. Wolfe, E. M. Gyorgy, P. L. Finn, M. H. Kryder, and C.-H. Chang, "Near-field magneto-optics and high density data storage," *Appl. Phys. Lett.* **61**, 142 (1992).
- ⁹² http://www.research.ibm.com/resources/magazine/1997/issue_2/storage297.html#one
- ⁹³ B.D. Terris, H.J. Mamin, M.E. Best, J.A. Logan and D. Rugar, S.A. Rishton, "Nanoscale replication for scanning probe data storage", *Appl. Phys. Lett.* **69**(27), p.4262, Dec. 1996.
- ⁹⁴ S. A. Miller, K. L. Turner, and N. C. MacDonald, "Microelectromechanical scanning probe instruments for array architectures," *Rev. Sci. Instrum.* **68**, p. 4155 (1997).
- ⁹⁵ G. L. Christenson, S. A. Miller, Z. H. Zhu, N. C. MacDonald, and Y. H. Lo, "Optical reading and writing on GaAs using an atomic force microscope," *Appl. Phys. Lett.* **66**, 2780 (1995).
- ⁹⁶ A. S. Dvornikov, I. Cokgor, F. B. McCormick, R. Piyaket, S. C. Esener, P. M. Rentzepis, "Molecular transformation as a means for 3D optical memory devices," *Optics Communications* **128** (pp. 205-210, (1996)
- ⁹⁷ F. B. McCormick, I. Cokgor, S. C. Esener, A. S. Dvornikov, and P. M. Rentzepis, "Two-photon absorption-based 3-D optical memories," in *High Density Data Recording and Retrieval Technologies*, Ted A Schwartz, Martin Francis, Editors, *Proc. SPIE* **2604**, 23-32 (1996).
- ⁹⁸ M. M. Wang, S. C. Esener, F. B. McCormick, I. Cokgor, A. S. Dvornikov, P. M. Rentzepis, "Experimental characterization of a two-photon memory," *Optics Letters* **22** (8), pp. 558-560, (1997).
- ⁹⁹ I. Cokgor, F. B. McCormick, A. S. Dvornikov, M. M. Wang, N. Kim, K. Coblenz, S. C. Esener, P. M. Rentzepis, "Multilayer disk recording using 2-photon absorption and the numerical modeling of the recording process," in *Optical Data Storage '97*, H. Bireki and J. Z. Kwiecien, Editors, *Proc. SPIE* Vol. **3109**, pp. 182-6, (1997).
- ¹⁰⁰ Lightwave Electronics Corp., Mountain View, CA.
- ¹⁰¹ ECR, San Diego, CA.
- ¹⁰² Microlase Optical Systems, Ltd, Glasgow, Scotland, UK.
- ¹⁰³ Sfim ODS, groupe sfim, Velizy Cedex, France.
- ¹⁰⁴ Time-Bandwidth Products, Inc., Zurich, Switzerland.

DISTRIBUTION LIST

MR. ALBERT JAMBERDINO AFRL/IFED 32 BROOKS RD. ROOM 321 ROME, NY 13441-4114	10
CALL/RECALL, INC 6160 LUSK BLVD. C-206 SAN DIEGO, CA 92121	5
AFRL/IFOIL TECHNICAL LIBRARY 26 ELECTRONIC PKY ROME NY 13441-4514	1
ATTENTION: DTIC-OCC DEFENSE TECHNICAL INFO CENTER 8725 JOHN J. KINGMAN ROAD, STE 0944 FT. BELVOIR, VA 22060-6218	2
ATTN: NAN PFRIMMER IIT RESEARCH INSTITUTE 201 MILL ST. ROME, NY 13440	1
AFIT ACADEMIC LIBRARY AFIT/LDR, 2950 P. STREET AREA B, BLDG 642 WRIGHT-PATTERSON AFB OH 45433-7765	1
AFRL/MLME 2977 P STREET, STE 6 WRIGHT-PATTERSON AFB OH 45433-7739	1
ATTN: SMDC IM PL US ARMY SPACE & MISSILE DEF CMD P.O. BOX 1500 HUNTSVILLE AL 35807-3801	1
COMMANDER, CODE 4TL000D TECHNICAL LIBRARY, NAWC-WD 1 ADMINISTRATION CIRCLE CHINA LAKE CA 93555-6100	1
CDR, US ARMY AVIATION & MISSILE CMD REDSTONE SCIENTIFIC INFORMATION CTR ATTN: AMSAM-RD-OB-R, (DOCUMENTS) REDSTONE ARSENAL AL 35898-5000	2

REPORT LIBRARY MS P364 LOS ALAMOS NATIONAL LABORATORY LOS ALAMOS NM 87545	1
AFIWC/MSY 102 HALL BLVD, STE 315 SAN ANTONIO TX 78243-7016	1
USAF/AIR FORCE RESEARCH LABORATORY AFRL/VSOSA(LIBRARY-BLDG 1103) 5 WRIGHT DRIVE HANSCOM AFB MA 01731-3004	1
ATTN: EILEEN LADUKE/D460 MITRE CORPORATION 202 BURLINGTON RD BEDFORD MA 01730	1
DOUSD(P)/DTSA/DUTD ATTN: PATRICK G. SULLIVAN, JR. 400 ARMY NAVY DRIVE SUITE 300 ARLINGTON VA 22202	1
MR. BERNARD J. CLARKE AFRL/IFED 32 BROOKS RD. ROOM 321 ROME, NY 13441-4114	1
MR. FRED HARITATOS AFRL/IFED 32 BROOKS RD. ROME, NY 13441-4114	1
DR. ED DANISZEWSKI AFRL/IFED 32 BROOKS RD, ROOM 321 ROME, NY 13441-4114	1
LT GARY SUNADA AFRL/IFED 32 BROOKS RD. ROOM 321 ROME, NY 13441-4114	1
DR ELIAS TOWE DARPA/MTD 3701 NORTH FAIRFAX DR. ARLINGTON, VA 22203	5

DR. ROBERT LEHANEY 1
DARPA/MTD
3701 NORTH FAIRFAX DR.
ARLINGTON, VA 22203

LTC DAVID HONEY 1
DARPA/MTD
3701 NORTH FAIRFAX DR
ARLINGTON, VA 22203

MR. GARY W. BARRINGER 1
AC2A/CCT
130 ANDREW ST. SUITE 216
LANGLEY AFB, VA 23665-1993

LTC STEVE PETERSEN 1
SAF/AQRT
1919 S. EADS ST. SUITE 100
ARLINGTON, VA 22202-3053

MIKE TINNEY 1
HQ AIA/XPD
#2 HALL BLVD. STE. 210
SAN ANTONIO, TX 78243-7010

MIKE POSIE 1
HQ AIA/XPD
2 HALL BLVD. STE. 210
SAN ANTONIO, TX 78243-7010

SUSAN BARMORE 1
HQ AIA/XPD
2 HALL BLVD. STE. 210
SAN ANTONIO, TX 78243-7010

GREGORY H. HART 1
DIA
IMAGERY EXPLOITATION AND INT. DIV.
PGX5B
WASHINGTON, DC 20340

***MISSION
OF
AFRL/INFORMATION DIRECTORATE (IF)***

The advancement and application of information systems science and technology for aerospace command and control and its transition to air, space, and ground systems to meet customer needs in the areas of Global Awareness, Dynamic Planning and Execution, and Global Information Exchange is the focus of this AFRL organization. The directorate's areas of investigation include a broad spectrum of information and fusion, communication, collaborative environment and modeling and simulation, defensive information warfare, and intelligent information systems technologies.